

Top 10 New Features in Windows 8.1

Windows IT Pro

A PENTON PUBLICATION

SEPTEMBER 2013 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Using Objects in PowerShell

Windows Server 2012
Failover Clustering

System Center 2012
Configuration Manager SP1

Windows Server 2012
Storage Live Migration

Windows Server 2012
Active Directory
Recycle Bin

FREE Newsletters!

**Not your average
Newsletters!**

WinInfo Daily UPDATE

Paul Thurrott covers the entire Windows universe with reviews, commentary, analysis, and tips. Delivered daily.

Windows IT Pro UPDATE

Windows industry news, products, FAQs, tips, and resources for IT professionals. Delivered weekly.

Cloud & Virtualization UPDATE

Get the latest news, blogs and analysis to help you determine your organization's cloud and virtualization strategy. Delivered weekly.

Exchange and Outlook UPDATE

News, strategies, products, and developments in Exchange Server and Outlook messaging. Delivered weekly.

Security UPDATE

Learn about Windows security risks, attacks, and how to fix or avoid them. Includes security alerts! Delivered bi-weekly.

Dev Pro UPDATE

Topics for Microsoft platform developers: ASP.NET, .NET Framework, Silverlight, mobile, and SQL Server development. Delivered weekly.

SQL Server Pro UPDATE

The latest news, products, and developments for SQL Server DBAs and developers. Delivered weekly.

SharePoint Pro UPDATE

SharePoint for IT professionals and developers – weekly tips, news, and how-to's. Delivered weekly.

**subscribe today at
windowsitpro.com/manage-newsletters**

WindowsITPro



Managing a Hybrid Virtualization Environment Online Conference | September 12, 2013

Intensive training and education all from the comfort of your own home or office!

Network and learn alongside your peers as industry experts provide in-depth presentations covering a wide range of virtualization topics.

Join Michael Otey, Alan Sugano and Orin Thomas as they uncover:

- Integrated Management for vSphere and Hyper-V
- Building the Hybrid Virtualized Cloud
- Monitoring a Mixed Hypervisor Environment with Operations Manager
- Best Practices for Virtualization Backup

You don't want to miss this free one-day event!

Click [here](#) for more information or to reserve your seat.

COVER STORY ▼

PowerShell Basics: Introduction to Objects

33

— Robert Sheldon

The more you understand how objects work within the PowerShell environment, the better you'll be able to take advantage of PowerShell's flexibility and perform the variety of tasks that PowerShell supports.



Features

46 Windows Server 2012: Building a Two-Node Failover Cluster

Michael Otey

73 Top 5 Features in System Center 2012 Configuration Manager SP1

Kent Agerlund

87 Windows Server 2012 Storage Live Migration

John Savill

Products

103 New & Improved

Interact

97 Ask the Experts

In Every Issue

107 Advertiser Directory

107 Directory of Services

107 Vendor Directory

Chat with Us



Facebook



Twitter



LinkedIn

Columns



6 [Need to Know](#)

Windows Server 2012 R2 Essentials and Windows Phone 8

Paul Thurrott



12 [Windows Power Tools](#)

Managing Printers Gets Easier in Windows 8

Mark Minasi



15 [Top 10](#)

Top 10 New Features in Windows 8.1

Michael Otey



18 [Enterprise Identity](#)

Bridging the Gap Between Active Directory and Office 365

Sean Deuby



23 [What Would Microsoft Support Do?](#)

Introducing the Active Directory Recycle Bin in Windows Server 2012

John Marlin

Editorial

Editorial Director: Megan Keller
Editor-in-Chief: Amy Eisenberg
Senior Technical Director: Michael Otey
Technical Director: Sean Deuby
Senior Technical Analyst: Paul Thurrott
IT Community Manager: Rod Trent
Systems Management, Networking,
Hardware: Jason Bovberg
Scripting: Blair Greenwood
SharePoint, Active Directory, Security,
Virtualization: Caroline Marwitz
SQL Server, Developer Content:
Megan Keller
Managing Editor: Lavon Peters
Editorial SEO Specialist: Jayleen Heft

Senior Contributing Editors

David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,
Jeff Felling, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas, Douglas Toombs,
Ethan Wilansky

Art & Production

Senior Graphic Designer: Matt Wiebe
Director of Production: Dylan Goodwin
Group Production Manager:
Julie Jantzer-Ward
Project Manager: Adriane Wineinger
Graphic Specialist: Karly Prickett

Advertising Sales

Technology Market Leader: Peg Miller
Key Account Director:
Chrissy Ferraro • 970-203-2883
Account Executives:
Megan Key • 970-203-2844
Barbara Ritter • 858-367-8058
Cass Schulz • 858-357-7649

Client Services

Senior Client Services Manager:
Michelle Andrews • 970-613-4964
Ad Production Coordinator: Kara Walby

Marketing & Circulation

Customer Service • 800-793-5697
Senior Director, Marketing Analytics:
Tricia Syed

Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

Corporate

Chief Executive Officer:
David Kieselstein
Chief Financial Officer/Executive Vice
President: Nicola Allais



List Rentals

MeritDirect
333 Westchester Avenue,
White Plains, NY 10604

Reprints

Reprint Sales:
Wright's Media • 877-652-5295

Windows IT Pro, September 2013, Issue No. 229,
ISSN 1552-3136. *Windows IT Pro* is published monthly by
Penton. Copyright ©2013 Penton. All rights reserved. No
part of this publication may be reproduced or distributed
in any way without the written consent of Penton.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525,
800-621-1544 or 970-663-4700. Customer Service:
800-793-5697.

We welcome your comments and suggestions about the
content of *Windows IT Pro*. We reserve the right to edit all
submissions. Letters should include your name and
address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can
submit articles to articles@windowsitpro.com.

Program Code: Unless otherwise noted, all programming
code in this issue is ©2013, Penton, all rights reserved.
These programs may not be reproduced or distributed
in any form without permission in writing from the
publisher. It is the reader's responsibility to ensure
procedures and techniques used from this publication are
accurate and appropriate for the user's installation. No
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®
are trademarks or registered trademarks of Microsoft
Corporation in the United States and/or other countries
and are used by Penton, under license from owner.
Windows IT Pro is an independent publication not
affiliated with Microsoft Corporation. Microsoft
Corporation is not responsible in any way for the editorial
policy or other contents of the publication.

Windows IT Pro

Windows Server 2012 R2 Essentials and Windows Phone 8



Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.

Email



Twitter



Website



In this month's column, I take a peek at two product lines that aren't discussed much in the context of the "Blue" wave of product updates. Those products are, of course, Windows Server Essentials and Windows Phone.

Windows Server 2012 R2 Essentials

Last month, I wrote about the advances in the [Preview version of the Windows 8.1 update](#) for [Windows 8](#) and Windows RT and noted that the combination of many small improvements resulted in a dramatically better experience than the initial releases of those OSs. But the Windows 8.1 Preview wasn't the only OS software that Microsoft debuted. The firm also released preview versions of [Windows Server 2012 R2](#), [System Center 2012 R2](#), and [SQL Server 2014](#), the first two of which should ship at roughly the same time as Windows 8.1 this fall. (SQL Server 2014 will ship early next year, I'm told.)

I'd like to focus on an offshoot of Windows Server 2012 R2 called Windows Server 2012 R2 Essentials, which is also now available in Preview form. You might recognize this product as the latest version of a release that was intended to replace Windows Small Business Server. But with the R2 release, it's become so much more than that.

The Essentials product line, as we now know it, debuted with Windows Server 2012, and as you might recall, Microsoft at the time had significantly simplified the Server family of products down to just four SKUs: Standard, Datacenter, Essentials, and Foundation. The first two, Standard and Datacenter, offer identical feature sets and differ only in pricing and virtualization rights: Datacenter supports

unlimited virtualized instances of the product on the same hardware, but Standard supports just two.

The other two are interesting and hit the low end of the market. Foundation is the entry-level product and supports just 15 users and offers no virtualization rights; the only way to acquire this version is with new low-end server hardware. Essentials, meanwhile, has bigger shoes to fill: It's aimed at small businesses with up to 25 users and has no virtualization rights. But it's a replacement for many previous products, including Windows Home Server, Windows Small Business Server (both Standard and Essentials), and Windows Storage Server Essentials.

In R2 guise, Essentials takes on a far greater role, moving up-market to support mid-sized businesses and integrating with key Microsoft online services such as [Office 365](#). It also supports more platform technologies, including those that are new to Windows Server 2012 R2 as well as unique support for Hyper-V based virtualization, as you'll see.

From a high level, Essentials R2 looks and works much like its predecessor, offering a friendly Dashboard admin console that most small businesses can use in lieu of the more complicated Windows Server tools, though the full set is available as well. The Dashboard features five main areas—Home, Users, Devices, Storage, and Applications—and more can be added by integrating with supported Microsoft online services or by installing add-ins.

As with the initial version of the software, the Dashboard assumes no admin expertise—indeed, the key target market of small businesses is understood to have no such personnel on site—and a clear set of Getting Started tasks such as *Add user account*, *Add server folder*, and *Set up Anywhere* access is presented up front so the server can be quickly configured.

But Essentials R2 also offers some new features. First, R2 can integrate with Office 365, Windows Azure Backup, Windows Intune, or an on-premises Exchange server out of the box, and Health Report functionality is built right in. In the initial Essentials 2012 release,

only Office 365, Azure Backup, and Health Report were available, and then only after the fact, by installing add-ins.

If you're familiar with Essentials 2012, you'll notice that the R2 version offers deeper integration with other platform features. If you integrate the local server with Office 365, you can now manage Exchange Online distribution groups, SharePoint Online libraries, and mobile devices directly from the Dashboard. Server folders support quotas, and you can now add folders from a second server in your environment to the server's shared folders. It integrates with BranchCache for the first time.

But there are even bigger changes under the hood. Essentials R2 can be installed directly to a physical server, as before, or to a virtual machine (VM) in Hyper-V Server on top of a physical installation of the server. (Previously, Essentials didn't even include Hyper-V support.) That latter installation type enables several Hyper-V-, Azure-, and externally hosted online services and features to offer deeper insight and integration capabilities with Essentials R2 and was a major feature request of server makers who want to provide their own value-added services on top of the product. Microsoft tells me that certain Azure services will require or work better with VM-hosted OSs and will thus open up new possibilities inside the Microsoft sphere as well. On-site, you can take advantage of Hyper-V's Live Migration and Replica functionality.

Essentials R2's move to the mid-sized business world is accomplished by a similarly intriguing change: With Windows Server 2012 R2 Standard and Datacenter, you can now install Essentials as a role called Windows Server Essentials Experience, providing all of the unique Essentials features—online services integration, simple PC and device management with centralized backup and File History integration, simple storage management, and more—to bigger businesses of up to 300 employees. So if you've started out with a small business version of Office 365 and Essentials, you can move up the stack and continue using the tools you're familiar with.

I'm currently using Essentials 2012 as the center of my own home-based business and will be migrating to Essentials R2 when the final version arrives later this year. In the meantime, I'll keep plugging away at the Preview release and see what else I can find. But so far, this is stacking up to be a pretty impressive release that improves on the core feature set, broadens the availability to larger businesses, and integrates with key online services. Not bad for an R2 release.

Windows Phone 8.1 and More

When Microsoft launched Windows Phone three years ago, it promised to continually update the product with new features and work with carriers to ensure that users wouldn't suffer from an Android-like uncertainty regarding the timing of these updates. Since then, of course, we've discovered that Windows Phone users are indeed second-class citizens in the eyes of the wireless carriers, and update delivery, always problematic, has slowed to a crawl. This behavior has affected Windows Phone 8 in obvious ways and plays a role in the delay of the Windows Phone 8.1 update, which one might have assumed would otherwise arrive alongside Windows 8.1 and the other many Microsoft product updates that will ship later this year.

So what happened? According to my sources, phone updates are complex and expensive to the carriers, so they like to do as few as possible, given Windows Phone's relatively small user base, and bundle several updates together for a single test and deploy phase. This reluctance has essentially pushed back Microsoft's planned Windows Phone 8 updates over time.

These include three sets of relatively minor Windows Phone 8 software updates called GDRs (or General Distribution Releases). GDR1, code-named Portico (back when the Windows Phone team was still using code names that ended in "o," a practice that has since mysteriously ceased), first arrived in very late 2012 and added some fixes for messaging, text replies to calls, Internet Explorer, Wi-Fi connectivity, and other features. Its messy delivery to different device types on

different carriers over several months ended up being a harbinger of events to come.

GDR2 was completed in April but didn't arrive until July, on two new phones: the Lumia 925 and Lumia 1020. At the time of this writing, it was expected to start shipping to existing handsets during July and August. This update adds support for CalDAV and CardDAV (now required for Gmail contacts and Google Calendar interoperability) and Data Sense and returns the FM radio back to compatible handsets. It also adds fixes for Xbox Music, Skype, Internet Explorer, the camera, and other features.

GDR3, originally expected in time for fall, will likely ship closer to the end of the year. Although I have no personal sources for this release, rumors suggest it will include support for so-called "phablet" handsets with 5" to 6" screens and 1080px (1920 × 1080) resolution. If true, that's most certainly the release that will serve at the center of Microsoft's (and Nokia's) Fall 2013 plans. In the previous three years, of course, Microsoft delivered a major new Windows Phone release each October.

It's not all bad news. Microsoft also revealed that it doubled the new 36-month support lifecycle for Windows Phone 8 software updates and doubled the previous lifecycle of 18 months, making it a lot friendlier for enterprises. But that schedule is also a hint that Windows Phone 8 won't be updated quickly. Previous versions of Windows Phone were replaced within a year. With Windows Phone 8, we could be looking at a long haul of mostly minor changes.

That said, looking to 2014, at least two major updates are coming. The only questions are timing and whether carriers will allow them through in a timely fashion. If there's anything worse than no software updates, it's knowing that updates exist that you can't get.

The first is an Enterprise Feature Pack that Microsoft plans to ship in the first half of 2014, adding features that its enterprise customers say are still missing from Windows Phone 8. These include S/MIME to sign and encrypt email; app-aware, auto-triggered VPN (as in

Windows 8.1 on PCs) for access to corporate resources behind the fire-wall; enterprise Wi-Fi support with EAP-TLS; enhanced Mobile Device Management (MDM) policies that will work across Intune as well as third-party device management solutions; and certificate management for user authentication enrollment, updating, and revocation.

And then of course there's Windows Phone 8.1. As with GDR3, I've not heard anything about this release directly, but rumors claim it will close the gap with Windows 8/RT/8.1 from functional and usability and SDK/API perspectives and will possibly include such changes as a notification center, better multitasking (with explicit app shut-down), and various changes to the built-in apps.

But it's pretty clear that Windows Phone 8.1 is still some time away, and it seems that the release keeps getting pushed back. Certainly, details remain murky.

There are further rumors of yet another Windows Phone "reset" for Windows Phone 9, which could very well be just another version of Windows but aimed at handsets. This makes plenty of sense to me given the improvements we see in Windows 8.1, which fully supports portrait mode, and in Windows RT guise in particular would make for a fine phone OS.

But let's not get ahead of ourselves. We don't even have GDR2 yet.

What's interesting is that Windows Phone 8 still stacks up really well against the iPhone and Android competition, especially Nokia's latest devices, such as the Lumia 1020, which sports an amazing DSLR-like 41-megapixel camera. But it's hard to overlook how great an advantage this platform could have if the wireless carriers would simply allow Microsoft to deliver updates to users. ■

Managing Printers Gets Easier in Windows 8

Use this excellent PowerShell cmdlet to manage your printer drivers



Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.

Email



Twitter



Website



No matter how fervently I wish printers would just go away, it seems they're here to stay—with all their attendant jams, empty trays, and constant hunger for toner and ink. (Wouldn't a *Print to Surface*, *Kindle*, *iPad*, or *Google Glass* feature be great?) Windows 8, however, offers some assistance in configuring and managing the infernal things in the form of 13 PowerShell cmdlets. This month, I'd like to introduce the first one—*add-printerdriver*—to you.

Amongst Windows 8 and Windows Server 2012's new PowerShell nouns are *printer*, *printerdriver*, and *printjob*. As we've seen many times over the years, what we want from printers (the actual hardware kind) is paper output, and that output comes from print jobs. Print jobs are created by the second meaning of the word *printer*—a purely software notion arising from when we build a connection between a given physical PC and a physical printer. To connect that physical PC to the physical printer, however, we need two other things: either a cable (probably USB) or a network connection between the two devices, and a printer driver. In other words, printer drivers enable the creation of *printers* (the software object type), printers create print jobs, and print jobs create the desired print output.

Step one, then, is to manage printer drivers. PowerShell's *get-printerdriver*, *add-printerdriver*, and *remove-printerdriver* do that. You must have local administrative powers to add or remove printer drivers with these cmdlets. (I point that out because, as you'll see, you don't need admin juice to add printers, once you have a given

printer's drivers loaded on a system.) The basic syntax for *add-printerdriver* is simple. For example, to load the driver for a Dell 1320C, you'd type

```
add-printerdriver "Dell Color Laser 1320c"
```

As you might guess, that driver name is a “magic name” that needs to exactly match a known driver; otherwise, you'll get errors (and unfortunately the cmdlet doesn't take wildcards). You can, however, uncover a driver's “magic name” fairly easily either by typing *get-printerdriver* on a system that already has that driver loaded or by poking around inside the .inf file associated with any printer driver. Browse the file a bit, and you'll see a block that looks like this:

```
; Model sections.
[DELL]
"Dell Color Laser 1320c" = DLHSNZP1,LPTENUM\DellColor_
    Laser_132010D5
"Dell Color Laser 1320c" = DLHSNZP1,USBPRINT\DellColor_
    Laser_132010D5
```

It goes on for a bit more, but you can see that there's a string after the vendor name, and that's usually the “magic name” for the driver.

Add-printerdriver is great, but it promises more than it delivers. The examples that you've seen so far show that you can fairly simply add new printer drivers to your system—as long as they're currently in the driver store. Of course, Windows 8 systems (like Windows 7 systems) ship with a lot of drivers, but what about installing new ones? According to the documentation, you can load a driver that isn't already in your system's driver store with the *-infpath* option:

```
add-printerdriver -name "Waxtronic 320" -infpath "c:\drivers
    \waxtronic\wi949.inf"
```


In fact, Microsoft seems so sure of this capability that the `-inpath` option and the driver name are built in as positional parameters, letting you type that cmdlet as

```
add-printerdriver "Waxtronic 320" "c:\drivers  
  \waxtronic\wi949.inf"
```

But, unfortunately, neither formulation seems to work, and even the most probing Bing or Google search—"Why doesn't the `-inpath` option work?"—brings either no help or clueless help. Next version, perhaps. Until then, the much older `Pnputil` command might help. First appearing in Windows Vista, `Pnputil` offers some nice driver management help, including the ability to install any given driver (or drivers, for that matter—it can tackle a whole folder of them) with the `-i` and `-a` switches, as in

```
pnputil -i -a "c:\drivers\waxtronic\wi949.inf"
```

This command usually works, and it will report back that it has installed that driver package as `OEMnn.inf` (where *nn* is a number) in the `C:\windows\inf` folder, letting you open up that `.inf` file with Notepad and search for its English driver name. Please understand, however, that no matter how many times you try `Pnputil` or *set-printerdriver*, you won't convince Windows 8 to use pre-Vista drivers.

Oh, and two more points here. First, when you're typing that printer driver's name in an *add-printerdriver* command, you needn't worry about case. Second, if you decide to give `Pnputil` a chance, remember that it's not PowerShell, and so it absolutely needs to see the `-i` and `-a` parameters in that order. Try typing *pnputil -a -i something*, and it'll just look at you confusedly and offer some examples of proper `Pnputil` syntax. OK, now that the drivers are in, let's use them—next month! ■

Top 10 New Features in Windows 8.1

Learn about the best new features in the forthcoming OS update

By now, it should come as no surprise that [Windows 8](#) has failed to take businesses by storm. According to [NetMarketShare](#), Windows 8 now has just a bit more market share than the much-reviled Windows Vista release. Microsoft hopes to turn around those negative perceptions with the latest release of [Windows 8.1](#) (formerly code-named Blue). Windows 8.1 addresses several of the major complaints that customers have had about Windows 8. The Windows 8.1 release is currently in preview (at press time), but at the Worldwide Partner Conference in Houston in early July, Microsoft announced that [Windows 8.1 would be available in August 2013](#). The Windows 8.1 release will be a free download from the Windows Store. Here are the top 10 new features in Windows 8.1



Michael Otey

is senior technical director for *Windows IT Pro* and *SQL Server Pro*.



Email

1 New Start Button

Answering the primary criticism of Windows 8's UI, Microsoft is returning the Start button to Windows 8.1. The new Start button is in the lower-left corner but unfortunately it doesn't work the way the Start button works in Windows 7 and earlier. Instead of showing your installed programs and administration options, the new Start button simply switches you over to the Metro Start screen. Right-clicking it or clicking and holding it displays a context menu, from which you can open Program and Features, Power Options, Event Viewer, System Device Manager, Network Connections, Task Manager, and a number of other handy options.

② Boot to Desktop

The ability to boot directly to the desktop addresses another complaint about Windows 8 and the new Start screen. With Windows 8.1, you can open the Taskbar and Navigation properties, then click Navigation and select the setting *Go to the desktop instead of Start when I sign in*. The next time Windows 8.1 starts, you'll boot into the desktop, allowing you to avoid the new Start screen altogether.

③ Different Tile Sizes

As with Windows Phone 8, the new Windows 8.1 Start screen offers a variety of tile sizes, including new large and small tile layouts. Using the different-sized tiles helps you more creatively customize and optimize your Windows 8.1 Start screen, and the new large tile size enables the tiles to display a lot more information.

④ New Apps Display

Another handy change in Windows 8.1 is the new Apps display, which you can access by swiping up on the Start screen or by clicking the down arrow. The Apps screen shows all the installed applications and lets you sort the new Apps display in several ways, including by name, date installed, and most used.

⑤ Work Folders

Another handy new feature in Windows 8.1 is the new Work Folders capability. Work Folders allows a user to sync data to his or her device from a folder located on a [Windows Server 2012 R2](#) server using the File and Storage Services role. Work Folders lets you keep local copies of your files on your devices, and the Work Folders feature will automatically synchronize your data to the central server.

⑥ Workplace Join

Workplace Join lets administrators grant access to some corporate resources for devices that aren't members of a Windows domain. If a

user registers his or her device, IT can grant the device limited access to corporate resources and enable some management policies for the device. This feature requires Server 2012 R2.

7 Deeper SkyDrive Integration

You get 7GB of free SkyDrive storage, and Windows 8.1 is able to use that as the default location for document storage as well as for backing up Windows 8.1 to SkyDrive. There are also options to sync your device settings, including the Start screen and the position of tiles and other desktop personalization settings. In addition, Microsoft will provide a new Windows RT SkyDrive app.

8 Side-by-Side Apps

Another annoying limitation of Windows 8 is the inability to display more than two apps at a time. Although Windows 8.1 still doesn't allow you to run apps in windows the way [Stardock's ModernMix](#) does, Windows 8.1's new Snap View feature lets you run up to eight apps side by side if you have two 2560 × 1600 monitors. Smaller 1280 × 780 displays are limited to running two apps side by side.

9 Auto-Triggered VPN

The VPN support in Windows 8.1 has been enhanced with improved support for a broader range of VPN clients. In addition, Microsoft is adding a new auto-trigger VPN capability to Windows 8.1 apps that enables apps that need to access the organization's resource through a VPN to automatically launch the VPN when the app starts.

10 New Windows Store

The Windows Store has been redesigned to be more like Google Play or Apple's App Store. There is a new Spotlight section to feature various apps, and there are lists of top free apps right on the home page. A new description page provides more information about each of the apps. Windows 8.1 apps will also now update automatically. ■

Bridging the Gap Between Active Directory and Office 365

Configuring SSO to Office 365 with native tools can be done—but third-party solutions do it better



Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.

Email



Twitter



When Microsoft reported its [third-quarter financial results in April](#), the company revealed that 25 percent of its enterprise customers have [Office 365](#). When you take trial subscriptions into account, the actual deployment numbers are almost certainly lower. Small-business deployments are probably higher because this service is particularly appealing to companies without a lot of pre-existing IT infrastructure. Nonetheless, a lot of companies are looking hard at Office 365—and they're probably running up against difficult single sign-on (SSO) requirements. This month, I want to take a look at the challenges of integrating Office 365 into your environment while preventing your users from encountering credentials prompts; I also want to show you how third-party identity management providers can make that connection easier.

Office 365 Integration Options

Office 365 can integrate with your existing on-premises environment in five ways, from essentially no integration to full SSO using your existing corporate credentials, as you see in Table 1. Three of these integration methods use your company—typically an Active Directory (AD) forest—as the identity provider. Two of them use the same user IDs and passwords for Office 365 as for your corporate directory. But only one provides seamless access to Office 365 as you'd access your local resources via SSO. For midsized-to-large enterprises, this method is usually the best way to go; any time you

Table 1: Five Ways to Integrate Office 365 with Your On-Premises Environment

Integration Options	Target Customer Segment	Scenario Supported	Directory Source of Authority	Hardware Requirements	IDP	User Logon Experience	Complexity
Portal	Small	Least	Cloud	No additional hardware required	Cloud	<ul style="list-style-type: none"> • Disjoint username and password • Enter credentials twice 	Low
PowerShell/ Directory GRAPH	Small/ Medium	Least	Cloud	No additional hardware required	Cloud	<ul style="list-style-type: none"> • Disjoint username and password • Enter credentials twice 	Medium
DirSync with Cloud Identities	Small/ Medium	Some limitations	On premises	Windows Server OS for DirSync appliance	Cloud	<ul style="list-style-type: none"> • Same username, disjoint password • Enter credentials twice 	Low
DirSync with Password Sync	Small/ Medium	Some limitations	On premises	Windows Server OS for DirSync appliance	Cloud	<ul style="list-style-type: none"> • Same username and password for on-premises and cloud • Enter credentials twice 	Low
DirSync with SSO	Medium/ Large	Most	On premises	<ul style="list-style-type: none"> • DirSync appliance • AD FS (or other STS) deployment 	On premises	<ul style="list-style-type: none"> • Same username and password for on-premises and cloud • Log on once if on premises 	High

require a password prompt, user confusion and support costs will increase.

Regardless of whether you're trying to connect to Office 365 or another cloud service provider (CSP) such as Google Apps, Internet SSO requires two major components. First, you can't have access to a cloud service without an identity on that service. Therefore, you must have a method to populate—then keep synchronized—your on-premises identities with the cloud service. Microsoft's utility for identity provisioning to Office 365 is the [Windows Azure Active Directory Synchronization Tool](#) (more efficiently known as DirSync). Third-party identity management tools also have their own account provisioning mechanisms, which I'll cover shortly.

Account Provisioning with DirSync

DirSync straightforwardly monitors and synchronizes local AD objects with Windows Azure AD. It's a one-way sync, which means your local AD objects are always authoritative over the synchronized objects in Azure AD. Of course, when you dig a little deeper, it's not quite so simple. (It never is.) DirSync will synchronize as many as 50,000 objects with no intervention; if you need to sync more than that, you need to call Microsoft support (to promise that you're not launching a Denial of Service—DoS—attack on Azure AD, I assume). Also if you need to sync more than 50,000 accounts, you must install a full instance of SQL Server 2008/SQL Server 2008 R2. The DirSync server must be a member of the forest in which it's syncing objects, and the server needs to be as tightly secured as a domain controller (DC), but it can't be installed on a DC. If you're using DirSync with the password synchronization option (which isn't necessary or recommended if you're using federation), password changes are replicated every two minutes, but other changes might take several hours. Fellow Directory Services MVP Sander Berkouwer wrote a blog post called [“Five Things you should know about using DirSync with Password Sync”](#) that describes some of these characteristics.

Federation with AD FS

Once you have your identities up in Office 365, to get SSO you need a way to authenticate them with where they came from—your company, as the identity provider. With identity federation, this authentication occurs through a component known as an identity bridge. Microsoft's general-purpose identity bridge is Active Directory Federation Services (AD FS), and third parties (notably IDaaS providers) have dedicated identity bridges.

Although it gets easier with each new release, designing and deploying a production AD FS installation isn't a trivial task. AD FS was previously not to be installed on a DC, but in Windows Server 2012 R2 the recommended configuration is to be installed on a DC. You should configure AD FS for high availability because your users

won't be able to log on to critical office functions without it. Thus, you need to set up a Windows Failover cluster with the AD FS role installed, deploy an AD FS proxy server in your corporate DMZ, and obtain and install public certificates. Once installed and in production, the AD FS installation must be monitored and updated, and you can't let your public certificates expire or your trusts will fail. Further, a Microsoft-only solution requires that your federated trust with Office 365 be with a single forest, so if you have accounts in multiple AD forests, or non-AD identity sources, you need to perform some kind of consolidation. Thomas Kemp, in "[Options for Federated Identity for Office 365, Part 2](#)," does a nice job summarizing some of the challenges in a local AD FS deployment.

Cleaning Up Dirty Directory Data

It doesn't matter how nicely you're connected to Office 365 if the data in your directory isn't clean. Different username and phone number conventions are common, for example, and although the Office 365 Deployment Readiness Tool will show you where your data problems are, correcting them and setting policies in place to keep them clean can take time. (Note that the Office 365 Deployment Readiness Tool, which used to be a free download that you could run against any AD forest to check for problems that would delay or prevent Office 365 migration, has been moved into the Office 365 subscription. To use it, you'll have to sign up for a trial subscription.)

This adds up to a solid set of challenges in the way of an Office 365 deployment. And it's not even Office 365 itself that's the challenge: It's the plumbing that's getting in the way.

Third-Party Solutions

As it does in so many scenarios, Microsoft provides basic functionality out of the box to get the job done, but it might not be the easiest or most fully featured implementation. Are you in a rush to deploy Office 365? Are you willing to spend some money to simplify the

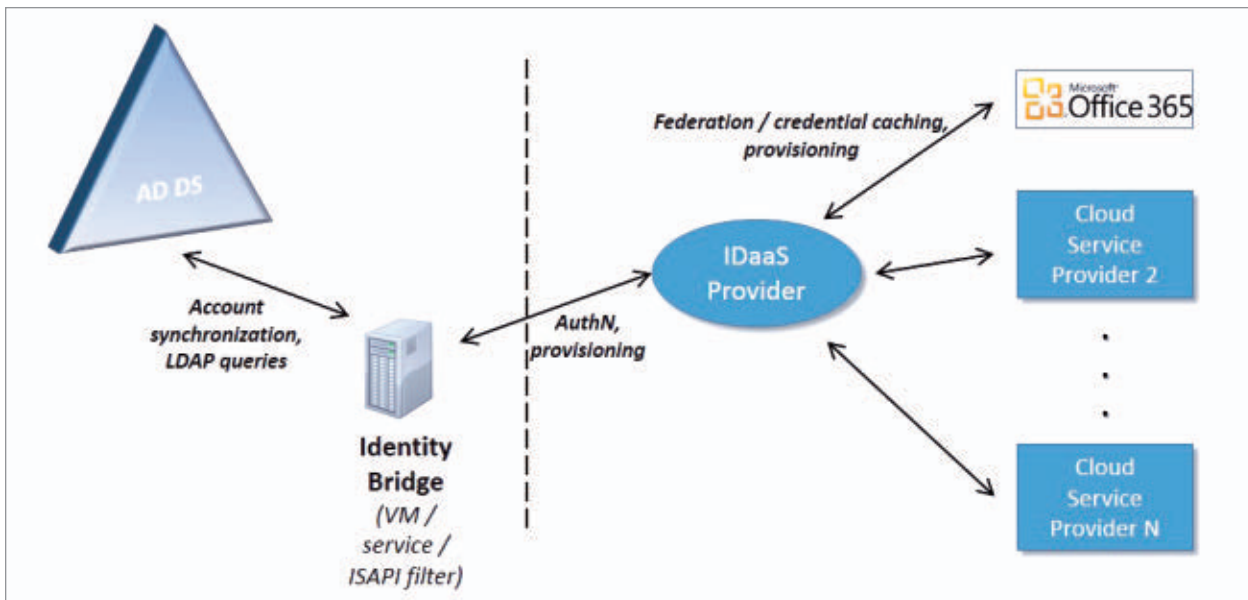


Figure 1
IDaaS Vendors/
Specialized Identity
Bridge

connection? Budgets are tight, but time is money, so in this case throwing cash at the problem really can help. Identity vendors will be happy to help ease your pain.

If your on-premises identity sources don't play well with one another, an on-premises virtual directory service (VDS) can make all of them appear to be a single directory. You can also create rules in the VDS to reformat your dirty data to present a clean view to offsite web services. If you don't want the hassle of hosting your own highly available AD FS cluster, you can find several on-premises identity bridges specifically designed to connect to Office 365 with a minimum of configuration work. And IDaaS vendors—with the easy-to-deploy, specialized identity bridge that Figure 1 shows—can quickly get you through the federation problem. They provide SSO capability to hundreds of cloud service providers, including Office 365. They also provide account provisioning capabilities, although their Office 365 account provisioning ability varies. Finally, many of these vendors provide all three of these important functions (federation, provisioning, and identity consolidation) within their product suites. ■

Introducing the Active Directory Recycle Bin in Windows Server 2012

Setting up and using this tool just got a whole lot easier

The Recycle Bin has been a feature in Windows OSs for quite some time. If you accidentally delete a file, it gives you the ability to restore it. However, if you accidentally delete a user or computer account in Windows Server 2008 Active Directory (AD) or earlier, you have very little you can do other than restore AD, re-create the account, or use a third-party tool to recover it.

Windows Server 2008 R2 introduced the first Active Directory Recycle Bin. With it, you can restore a user, computer, or organizational unit (OU) account that has been accidentally deleted. However, you must use Windows PowerShell to work with the Active Directory Recycle Bin. Using PowerShell commands to search for and restore a deleted object can be difficult, especially if you're not familiar with PowerShell. And when you want to search for an object, you're limited in what you can search on.

The PowerShell commands can also get a bit on the long side. For example, here's the command to enable the Active Directory Recycle Bin feature:

```
Enable-ADOptionalFeature "Recycle Bin Feature" -server `
((Get-ADForest -Current LocalComputer).DomainNamingMaster) `
-Scope ForestOrConfigurationSet `
-Target (Get-ADForest -Current LocalComputer)
```



**John
Marlin**

is a senior support escalation engineer in Windows Commercial Technical Support, focusing on failover clustering. He is a Microsoft Certified Trainer for clustering, delivering to Microsoft and its partners, and is a regular contributor to the [Ask the Core Team](#) blog. He is also a contributor to the new book *Introducing Windows Server 2012* (Microsoft Press).



Email



Website

If you want to search through all the deleted objects in the Active Directory Recycle Bin, you need to run the command:

```
Get-ADObject -filter 'isDeleted -eq $true' `
-and name -ne "Deleted Objects" -includeDeletedObjects `
-property * |
ft msds-lastKnownRdn,lastKnownParent -auto -wrap
```

If you want to restore a user named JohnMarlin, your command would be:

```
Get-ADObject -Filter 'samaccountname -eq "JohnMarlin"' `
-IncludeDeletedObjects | Restore-ADObject
```

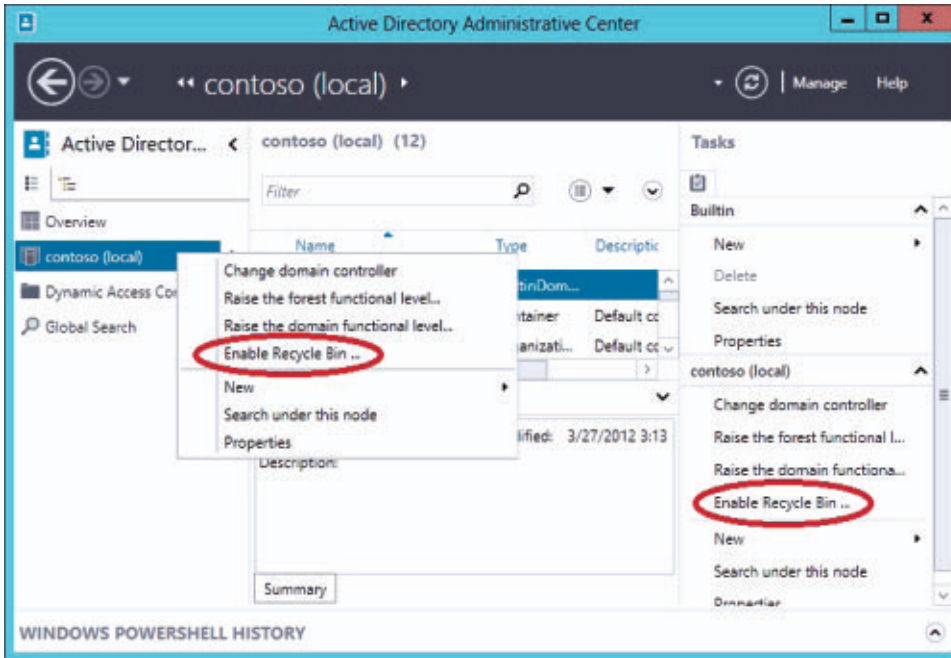
As you can see, these aren't simple commands and the chances for typos are there. Don't get me wrong—I'm not saying the Active Directory Recycle Bin in Server 2008 R2 is a bad thing. It's just that the way to manipulate it tends to be challenging.

These challenges were presented to the appropriate Microsoft Product Groups. Based on this feedback, they made the Active Directory Recycle Bin part of the Active Directory Administrative Center in Windows Server 2012. As you'll see, setting up and using the Active Directory Recycle Bin just got a whole lot easier.

You can restore not only a single object but also multiple objects simultaneously.

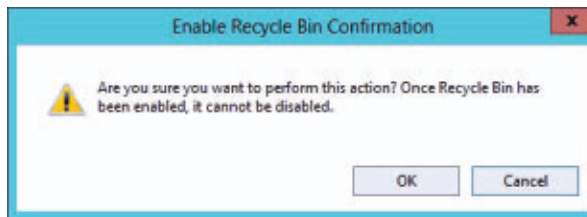
Setting Up the Active Directory Recycle Bin

Like its predecessor, the Active Directory Recycle Bin in Server 2012 isn't enabled by default and requires a Server 2008 R2 or later Forest Functional Level. To enable the recycle bin in Server 2012, you need to open the Active Directory Administrative Center, click your domain's name, and select Enable Recycle Bin from the Tasks menu. Alternatively, you can right-click your domain's name and select Enable Recycle Bin from the context menu. Figure 1 shows both methods.

**Figure 1**

Selecting the Option to Enable the Recycle Bin

After you select the Enable Recycle Bin option, you'll receive the Enable Recycle Bin Confirmation dialog box shown in Figure 2. As it notes, once you enable the recycle bin, it will always be enabled. You can't disable it at a later time.

**Figure 2**

Confirming That You Want to Enable the Recycle Bin

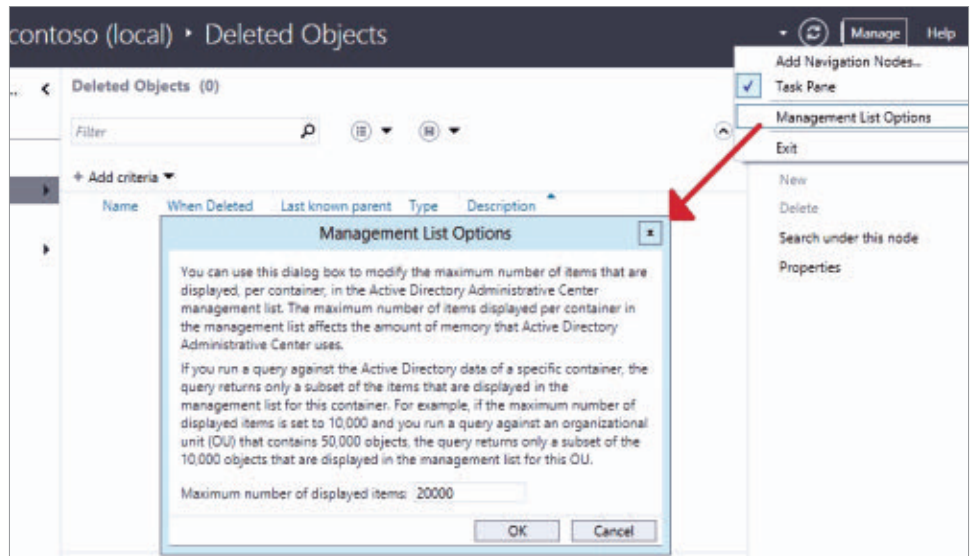
Before you enable the recycle bin, you also need to be aware that the size of the AD database (Ntds.dit) will increase. The disk space used by the recycle bin will continue to increase over time as it preserves objects and their attribute data. So, you need to make sure that you won't run out of disk space, especially if you're in the habit of continually deleting objects from AD. Note that you must be a member of the Enterprise Administrators Group to access the recycle bin.

After you click OK in the Enable Recycle Bin Confirmation dialog box, you'll be reminded that the recycle bin won't be fully functional until the change is replicated to all remaining domain controllers (DCs). Once fully functional, when you delete an object, it's saved based on the information in the msDS-deletedObjectLifetime attribute. This attribute describes how long a deleted object will be restorable. It's set in the CN = Directory Service,CN = Windows NT,CN = Services, CN = Configuration,DC = COMPANY,DC = COM container.

By default, the msDS-deletedObjectLifetime attribute is set to match the forest's tombstoneLifetime attribute. This attribute has existed since Windows 2000. Although its default had been 60 days, it was increased to 180 days in Windows Server 2003 SP1, where it continues to remain the default. The tombstoneLifetime attribute is set in the CN = Directory Service,CN = Windows NT,CN = Services, CN = Configuration,DC = COMPANY,DC = COM container.

There's an artificial limit to the number of items displayed in the recycle bin. By default, it's set at 20,000. You can change this number (up to 100,000) by clicking the Manage menu and selecting Management List Options, as shown in Figure 3.

Figure 3
Changing the Number
of Items Displayed in
the Recycle Bin



Using the Active Directory Recycle Bin

To access the Active Directory Recycle Bin in the Active Directory Administrative Center, you need to select the Deleted Objects option that appears under your domain name. By default, there will be five columns: Name, When Deleted, *Last known parent*, Type, and Description. You can add columns by right-clicking a column to get the additional options, as Figure 4 shows.

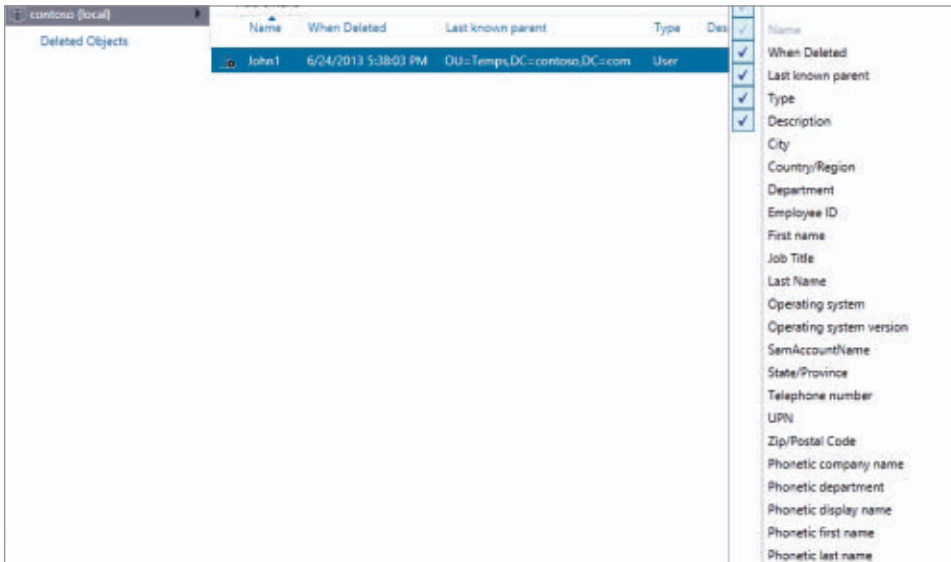


Figure 4
Adding Columns

After the recycle bin has been running for a while, you'll likely have several thousand objects in it. Scrolling through the list will take quite a long time. Fortunately, you can search with filters to narrow down what you need to find. For example, suppose that a temporary user account that has John in the name was deleted by mistake and you need to restore it. You're not completely sure which user account you need to restore because you didn't delete it, but you have some information about this user:

- He works in the Dallas, Texas, office.
- He's part of the accounting department.
- He hasn't logged on for 10 days.

Figure 5
Reviewing the Criteria
Available to Narrow a
Search

+ Add criteria ▼

- ☐ Users with disabled/enabled accounts.
- ☐ Users with an expired password.
- ☐ Users whose password has an expiration date/no expiration date.
- ☐ Users with enabled but locked accounts.
- ☐ Users with enabled accounts who have not logged on for more than a given number of days.
- ☐ Users with a password expiring in a given number of days.
- ☐ Computers running as a given domain controller type.
- ☐ Last modified between given dates.
- ☐ Object type is user/inetOrgPerson/computer/group/organization unit.
- ☐ Directly applied password settings for a specific user.
- ☐ Directly applied password settings for a specific global security group.
- ☐ Resultant password settings for a specific user.
- ☐ Resource property lists contain a given resource property.
- ☐ Name
- ☐ When Deleted
- ☐ Last known parent
- ☐ Type
- ☐ Description
- ☐ City
- ☐ Country/Region
- ☐ Department
- ☐ Employee ID
- ☐ First name
- ☐ Job Title
- ☐ Last Name
- ☐ Operating system
- ☐ Operating system version
- ☐ SamAccountName
- ☐ State/Province
- ☐ Telephone number
- ☐ UPN
- ☐ Zip/Postal Code
- ☐ Phonetic company name
- ☐ Phonetic department
- ☐ Phonetic display name
- ☐ Phonetic first name
- ☐ Phonetic last name

Figure 6
Adding the Criteria
Needed to Find John's
User Account

John

and City starts with Dallas

and Users with enabled accounts who have not logged on for more than this number of days: 15

and Users with an expired password.

and Department starts with Accounting

+ Add criteria ▼

Name	When Deleted	Last known parent	Type	Description
John23	6/24/2013 6:18:55 PM	OU=Temps,DC=contoso,DC=com	User	

- When he last logged on, he had 2 days before his password expired.

If you type John in the Filter box, it'll show everything with the name John in it. But what if there are hundreds of accounts with John in the name? You can add criteria to narrow the search by clicking the *Add criteria* button. As Figure 5 shows, there are many criteria from which to choose. Based on what you know about the user, you add several criteria and enter the specific information, as shown in Figure 6. As you can see, John23 is the user account you need to restore.

To restore the John23 object, you can right-click it and select the Restore option

(restores it to the original OU) or the Restore To option (restores it to another OU that you select). These two options are also available from the Tasks menu.

You can restore not only a single object but also multiple objects simultaneously. You can even restore an OU. For example, suppose you work for a company that sometimes uses temporary employees. There's a group of temporary employees whose contracts will end on Friday afternoon. As the head administrator, you're responsible for removing the OU (Temp-Employees) and all the user accounts in it when those contracts are done. On Thursday afternoon, you decide to take off Friday since you'll be on vacation the following week. To make sure that everything is taken care of, you write a script that will run late Friday night to delete those objects.

On Friday afternoon, upper management decides to keep the temporary employees on for another week to complete what they're doing. You were sent an email telling you to delay deleting the objects for a week, but you weren't there to receive it. On Friday night, your script runs and deletes the Temp-Employees OU as well as all the user accounts in it. Monday morning comes and the temporary employees can't log on. One of the other administrators, Mike, goes to the recycle bin to restore the user accounts. However, he's unaware that the OU was deleted. He also doesn't know the names of the users.

After opening the recycle bin, Mike adds the criteria shown in Figure 7 and gets the list of user accounts. However, when he tries to restore the first user account, he gets the pop-up box that Figure 8 shows. From the error message, Mike determines that the Temp-Employees OU has also been deleted. The recycle bin won't allow you to restore to an OU that doesn't exist, and it won't create one for you. So, Mike runs a new search to find the deleted OU and restores it. He then reruns the previous search to bring up the list of deleted user accounts, selects all of them, and restores them in one simple operation, as shown in Figure 9. Now all the temporary employees can log on again.

Figure 7

Adding the Criteria Needed to Find the Temporary Employees' User Accounts

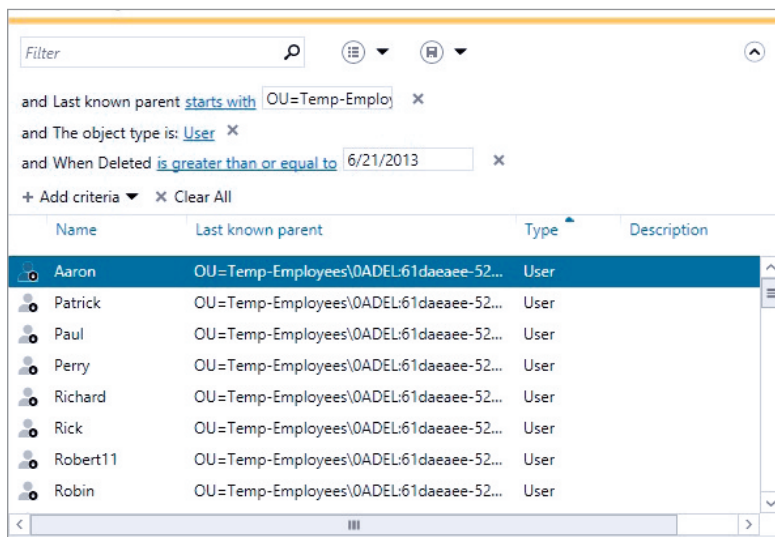


Figure 8

Receiving an Error Message

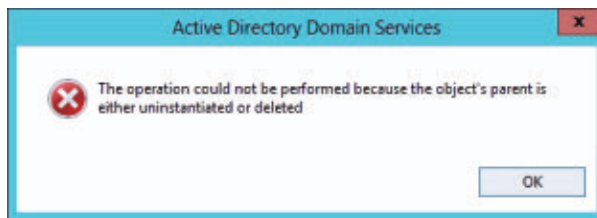
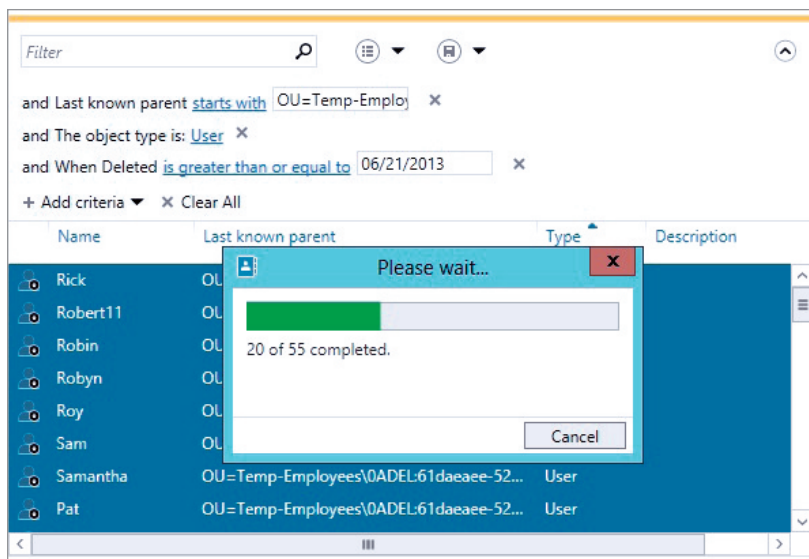


Figure 9

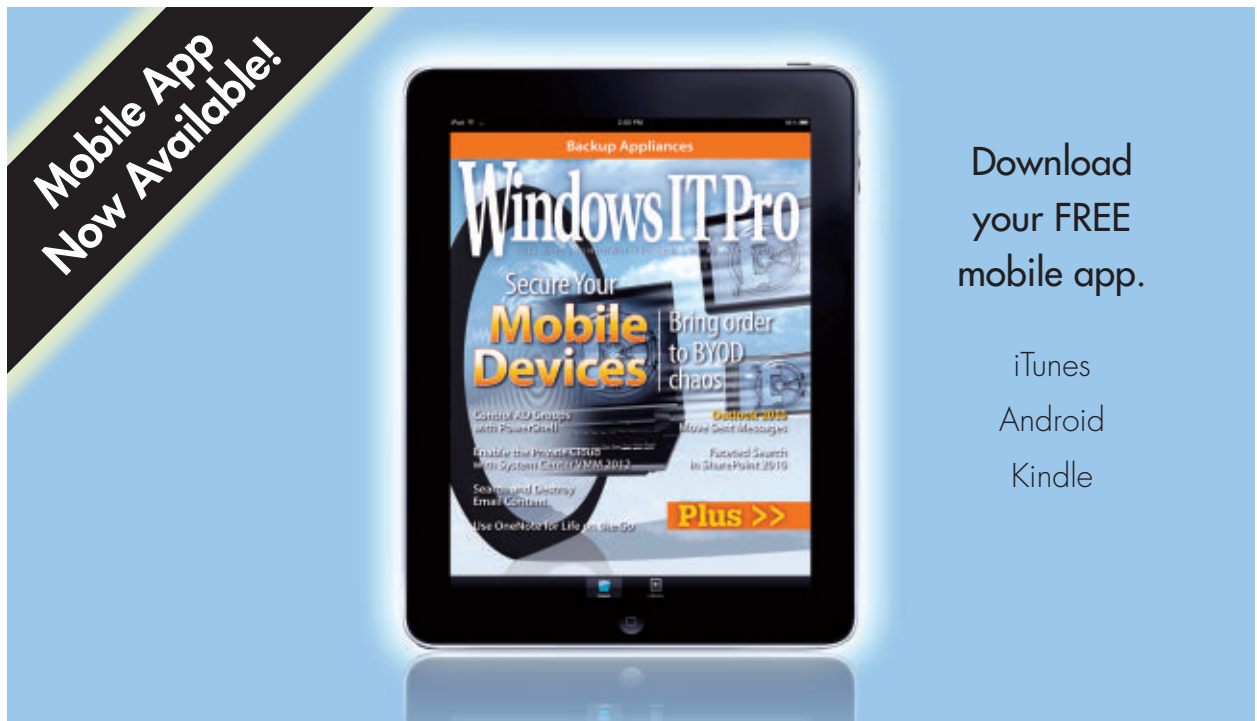
Restoring All the User Accounts at Once



As you probably are aware, AD has multiple partitions. It's important to note that the recycle bin can manage only domain partitions. So, if objects are deleted from the Configuration, Domain DNS, or Forest DNS partitions, you can't restore them with this tool.

A Lifesaver

The Active Directory Recycle Bin can be a lifesaver for those times when simply re-creating a user won't do or when you have to restore the entire AD or large parts of it. If the need arises, I encourage you to take advantage of this easy-to-use tool. ■



**Mobile App
Now Available!**

Backup Appliances

Windows IT Pro

Secure Your Mobile Devices

Bring order to BYOD chaos

October 2013

Plus >>

Download your **FREE** mobile app.

iTunes
Android
Kindle



Windows IT Pro Store

eLearning Classes

eBooks

On-Demand Training

In-Person Training

Posters

Videos

Plus you can **RENEW** your subscription or
UPGRADE to VIP membership while
you're there!

Stop by the store today!

Windows® IT Pro

PowerShell Basics

Introduction to Objects

Understanding and working with PowerShell objects

The ability to use [Windows PowerShell](#) effectively relies on an understanding of how objects are utilized to move and manage data as it passes through the PowerShell pipeline. The pipeline provides a structure for creating complex scripts that are broken down into one or more simple commands, each performing a discrete action against the data as it passes through. Objects make it possible to hand off that data from one command to the next by bundling it into individual packages of related information.

PowerShell objects provide a consistent structure for working with different types of data, regardless of that data's source. In other words, the ways in which you manage one object's data are similar to those used to manage another object's data. Because of PowerShell's object-oriented nature, you're able to not only take advantage of the objects inherently generated by PowerShell's built-in cmdlets but also build your own objects based on classes in the Microsoft .NET Framework. Indeed, PowerShell's use of objects—and the flexibility they provide—is what makes PowerShell such an effective and formidable tool.



Robert Sheldon

has written numerous books and articles about Windows technologies, database systems, business intelligence, scripting, enterprise operations, and consumerization. His books include *Beginning MySQL* (Wiley) and *SQL: A Beginner's Guide* (McGraw-Hill).



Email



Website

Understanding PowerShell's Object-Oriented Structure

The .NET Framework is a software-based structure that includes a large library of different types of classes. These classes serve as the foundation on which .NET objects are built and provide access to a variety of system, network, directory, and storage resources. PowerShell is built on specialized .NET classes that make it possible to access the entire .NET class library from within the PowerShell environment. You could say that objects provide the foundation on which PowerShell is built, which accounts for PowerShell's flexible and powerful nature.

Built into the PowerShell environment is a collection of cmdlets. Each cmdlet carries out a specific operation, such as retrieving a list of files in a folder or managing a service running on a computer. To carry out such an operation, the cmdlet generates an object or set of objects based on the specialized PowerShell classes. Objects provide the vehicles by which data is passed down the pipeline, where it can be used by other commands.

You can think of each object as a package of related information. For example, an object might contain the data necessary to describe a file: its name, size, location, and other attributes. To work with an object's data, you call its *members*, which are components that let you access and manipulate that information. A PowerShell object supports several types of members, but the two most common are *properties* and *methods*. A property is a named data value that describes the "thing" being represented by the object, such as the size of a file or the date it was created. Methods are actions that you can take related to the object's data, such as deleting or moving a file.

Working with PowerShell Objects

To better understand how PowerShell objects work, let's look at an example. One of the cmdlets that PowerShell supports, `Get-Service`, returns a list of services installed on a computer. When specified without any parameters, the cmdlet provides details about the services on the local computer.

Each service returned by the `Get-Service` cmdlet is an object based on the .NET class `System.ServiceProcess.ServiceController`. Like most objects, `ServiceController` contains numerous properties and methods that you can use to access the data contained within that object or run operations against the data. For example, the `ServiceController` object includes the `Name` and `DisplayName` properties. The data value associated with the `Name` property provides the service's actual name. The data value associated with the `DisplayName` property provides the display name used for that service. The `ServiceController` object also includes a number of methods. For instance, you can use the `Start` method to launch the service represented by the object or use the `Stop` method to stop that service.

You're not expected to know every property, method, or other member type associated with an object. PowerShell provides the `Get-Member` cmdlet, a handy tool for accessing details about each member supported by an object. You can use this cmdlet to retrieve information about all the members or certain member types. For details about all the members, you first specify the cmdlet you want to learn about and then specify the `Get-Member` cmdlet. You need to separate the two cmdlets with a vertical bar (pipe), like this:

```
Get-Service | Get-Member
```

In this case, the command is piping the objects returned by the `Get-Service` cmdlet to the `Get-Member` cmdlet. When you run the command, you receive the results shown in Figure 1. As you can see, the list contains each member's name, member type, and definition, which might or might not make sense, depending on your programming background. In any case, you can see that the `ServiceController` object supports a number of members, mostly methods and properties. (You can find details about the different member types in the MSDN article "[PSMemberTypes Enumeration](#).")

Figure 1

Getting Information
About All the
Members of the
Get-Service Cmdlet

TypeName: System.ServiceProcess.ServiceController		
Name	MemberType	Definition
Name	AliasProperty	Name = ServiceName
RequiredServices	AliasProperty	RequiredServices = ServicesDependedOn
Disposed	Event	System.EventHandler Disposed(System.Object ...
Close	Method	void Close()
Continue	Method	void Continue()
CreateObjRef	Method	System.Runtime.Remoting.ObjRef ...
Dispose	Method	void Dispose(), void IDisposable.Dispose()
Equals	Method	bool Equals(System.Object obj)
ExecuteCommand	Method	void ExecuteCommand(int command)
GetHashCode	Method	int GetHashCode()
GetLifetimeService	Method	System.Object GetLifetimeService()
GetType	Method	type GetType()
InitializeLifetimeService	Method	System.Object InitializeLifetimeService()
Pause	Method	void Pause()
Refresh	Method	void Refresh()
Start	Method	void Start(), void Start(string[] args)
Stop	Method	void Stop()
WaitForStatus	Method	void WaitForStatus(System.ServiceProcess. ...
CanPauseAndContinue	Property	bool CanPauseAndContinue {get;}
CanShutdown	Property	bool CanShutdown {get;}
CanStop	Property	bool CanStop {get;}
Container	Property	System.ComponentModel.IContainer ...
DependentServices	Property	System.ServiceProcess.ServiceController[]...
DisplayName	Property	string DisplayName {get;set;}
MachineName	Property	string MachineName {get;set;}
ServiceHandle	Property	System.Runtime.InteropServices.SafeHandle ...
ServiceName	Property	string ServiceName {get;set;}
ServicesDependedOn	Property	System.ServiceProcess.ServiceController[] ...
ServiceType	Property	System.ServiceProcess.ServiceType...
Site	Property	System.ComponentModel.ISite Site {get;set;}
Status	Property	System.ServiceProcess. ...
ToString	ScriptMethod	System.Object ToString();

Notice that the results returned by the Get-Member cmdlet begin by listing the type name of the class on which the object is based—in this case, System.ServiceProcess.ServiceController. Having this information can be useful if you're familiar with .NET classes and want to better understand the operations you're trying to perform. You can also use it to make sure that you're working with the type of object with which you want to be working. This leads to another aspect of the information returned by the Get-Member cmdlet. As you might have noticed, the results include details about only a single object, even though the Get-Service cmdlet returns an object for each service. When the Get-Member cmdlet sees that multiple objects of the same type are being returned, it includes only one instance of those objects to avoid redundancy. Another way to look at this is to say that it returns only the class on which all these objects are based.

However, if a command returns more than one type of object, Get-Member returns information on each of those types. For example, if

you run the `Get-ChildItem` cmdlet against a directory that contains both files and folders, `Get-Member` will return details about the `System.IO.DirectoryInfo` class and the `System.IO.FileInfo` class. However, if the directory contains only files, `Get-Member` returns details about `System.IO.FileInfo` only. So, when you use `Get-Member`, you should make sure that you're viewing the object types you want to be viewing.

The `Get-Member` cmdlet also lets you view a list made up of a specific member type, such as properties or methods. To do so, you need to include the `-MemberType` parameter, followed by the type name. For example, the following command specifies that only the object's properties be returned:

```
Get-Service | Get-Member -MemberType Property
```

Being able to retrieve the details about an object's members—whether all members or a specific type of member—makes it easier to access specific information within that object. For example, you now know that the `ServiceController` class supports the `Status` and `ServiceType` properties. You can use that information within your commands to refine your operations. For instance, the following command uses these properties with the `Where-Object` cmdlet to filter the results:

```
Get-Service | Where-Object {$_.Status -eq "Running" `
    -and $_.ServiceType -eq "Win32OwnProcess"}
```

In this command, the `Get-Service` results (i.e., the set of `ServiceController` objects) are piped to the `Where-Object` cmdlet. For each object passed to the `Where-Object` cmdlet, you can access that object's properties in order to create the filters. To do so, you first specify the `$_` symbol—a system variable that points to the current object in the pipeline—followed by a period and the name of the property. You can then use those properties to return specific data by defining Boolean expressions that evaluate to either true or false. In this case,

you first specify that the Status property must equal *Running* and the ServiceType property must equal *Win32OwnProcess* in order for an object's data to be returned. Notice that the *-eq* comparison operator is used for the equal operator and the *-and* operator is used to join the two expressions. As a result, both conditions must evaluate to true in order for the object to be returned. Also notice the back tick (`) on the first line. This indicates that the code continues on the next line.

Figure 2 shows the results returned by the command. These are the services running on my system that meet the Where-Object requirements—that is, the services are running and have a service type of Win32OwnProcess. (Note that I'm running PowerShell in a Windows 7 virtual environment.)

Figure 2
Retrieving a
List of Filtered
ServiceController
Objects

Status	Name	DisplayName
Running	AdobeActiveFile...	Adobe Active File Monitor V6
Running	COMSysApp	COM+ System Application
Running	DefaultTabUpdate	DefaultTabUpdate
Running	MSDTC	Distributed Transaction Coordinator
Running	MSSQL\$SQLSRV2012	SQL Server (SQLSRV2012)
Running	MSSQLFDLauncher...	SQL Full-text Filter Daemon Launche...
Running	osppsvc	Office Software Protection Platform
Running	Parallels Coher...	Parallels Coherence Service
Running	Parallels Tools...	Parallels Tools Service
Running	Pml Driver HPZ12	Pml Driver HPZ12
Running	PrLvsProvider	PrLvsProvider
Running	SQLBrowser	SQL Server Browser
Running	SQLWriter	SQL Server VSS Writer
Running	VisualSVNServer	VisualSVN Server
Running	WSearch	Windows Search
Running	YahooAUService	Yahoo! Updater

Knowing the names of the properties supported by the Service Controller object makes it much easier to find the exact information for which you're looking. You can take this a step further by tagging on the Format-Table cmdlet so that the results are more readable:

```
Get-Service | Where-Object {$_.Status -eq "Running" `
    -and $_.ServiceType -eq "Win32OwnProcess"} |
Format-Table -AutoSize
```

Although the original example had already output the information in a tabular format, the Format-Table cmdlet lets you include the

Status	Name	DisplayName
Running	AdobeActiveFileMonitor6.0	Adobe Active File Monitor V6
Running	COMSysApp	COM+ System Application
Running	DefaultTabUpdate	DefaultTabUpdate
Running	MSDTC	Distributed Transaction Coordinator
Running	MSSQL\$SQLSRV2012	SQL Server (SQLSRV2012)
Running	MSSQLFDLauncher\$SQLSRV2012	SQL Full-text Filter Daemon Launcher (SQLSRV2012)
Running	osppsvc	Office Software Protection Platform
Running	Parallels Coherence Service	Parallels Coherence Service
Running	Parallels Tools Service	Parallels Tools Service
Running	Pml Driver HPZ12	Pml Driver HPZ12
Running	PrIVssProvider	PrIVssProvider
Running	SQLBrowser	SQL Server Browser
Running	SQLWriter	SQL Server VSS Writer
Running	VisualSVNServer	VisualSVN Server
Running	WSearch	Windows Search
Running	YahooAUService	Yahoo! Updater

Figure 3

Preventing Truncation with the Format-Table Cmdlet's -AutoSize Parameter

-AutoSize parameter, which prevents data from being truncated, as shown in Figure 3.

Now let's explore another advantage of being able to view an object's members. Let's start by returning information about a specific service, MsDtsServer110, which is the service used to run SQL Server Integration Services (SSIS). To get its information, you use the Get-Service cmdlet's -Name parameter to specify the service name:

```
Get-Service -Name MsDtsServer110 | Format-List
```

By piping the Get-Service cmdlet's results to the Format-List cmdlet, you can easily see the various property values associated with this ServiceController object, as shown in Figure 4.

```
Name           : MsDtsServer110
DisplayName     : SQL Server Integration Services 11.0
Status         : Stopped
DependentServices : {}
ServicesDependedOn : {}
CanPauseAndContinue : False
CanShutdown    : False
CanStop        : False
ServiceType    : Win32OwnProcess
```

Figure 4

Getting the Property Values Associated with the MsDtsServer110 Service

PowerShell also lets you access a specific property value within an object. For instance, suppose you only need to know the value of the Status property associated with the MsDtsServer110 service. One way

you can do this is to specify the Format-Wide cmdlet along with its -Property parameter:

```
Get-Service -Name MsDtsServer110 |  
    Format-Wide -Property Status
```

In this case, the command returns only the value *Stopped*, which is the value associated with the service's Status property. Alternatively, you can access the ServiceController object and its property directly:

```
(Get-Service -Name MsDtsServer110).Status
```

Notice that the Get-Service command is enclosed in parentheses. This is necessary because, without them, PowerShell would treat *MsDtsServer110.Status* as the complete name of the service rather than the service name followed by the Status property. The parentheses force PowerShell to first execute the command, which generates a ServiceController object. PowerShell then retrieves the Status property from that object. Once again, the command returns the value *Stopped*.

Another way you can achieve the same result is to assign the outputted ServiceController object to a variable, then use that variable to call the property value:

```
$ssis = Get-Service -Name MsDtsServer110  
$ssis.Status
```

As you can see, the \$ssis variable is defined and the output of the Get-Service command is assigned as its value. In this case, you don't have to enclose the Get-Service command in parentheses. No properties are being called at this point, so there's no confusion as to what you're assigning to the variable. The command's job is

to produce a `ServiceController` object, and that object becomes the variable's value. You can then use the `$ssis` variable to reference the `ServiceController` object's `Status` property value by adding a period and the property name. Like the two other commands, this command returns the value *Stopped*.

You're not limited to only the `Status` property. Using the `$ssis` variable, you can access any of the other property values. For instance, the following command uses the `$ssis` variable to access the `DisplayName` property:

```
$ssis.DisplayName
```

This command returns the value *SQL Server Integration Services 11.0*.

It's important to note that assigning an object to a variable this way locks the data at that point in time. For example, if the service's status should change, the variable's `Status` property wouldn't reflect the changed state because the original data had been assigned to that object.

Even so, this doesn't prevent you from using the variable's methods to take actions against the service. But before you do that, let's view a list of the available methods. To do so, pipe the `$ssis` variable's contents to the `Get-Member` cmdlet, specifying `Method` as the member type:

```
$ssis | Get-Member -MemberType Method
```

Not surprisingly, the command returns a list of methods associated with the `ServiceController` object—in this case, the object generated for the `MsDtsServer110` service. You can then use the `$ssis` variable to call any of these methods. For example, the following command starts the `MsDtsServer110` service:

```
$ssis.Start()
```

Objects provide the vehicles by which data is passed down the pipeline, where it can be used by other commands.

Note that when you call a method, you have to include the parentheses at the end, even if you're not passing any parameters to the method. Once the service has started, you can just as easily stop it:

```
$ssis.Stop()
```

As you can see, understanding how to work with the objects that the cmdlets generate can help you use PowerShell more effectively. However, PowerShell doesn't provide a cmdlet for every .NET class, yet there might be times when you want to use one of those classes to perform a particular operation. In such cases, you can create your own objects based on those classes.

Creating a .NET Object in PowerShell

The .NET Framework and its class library provide a comprehensive set of classes for carrying out a wide range of operations. Although PowerShell's built-in cmdlets go a long way in accessing those classes, not all relevant operations are addressed. For that reason, PowerShell includes the New-Object cmdlet so that you can create your own objects.

It's important to note that in addition to creating objects based on .NET classes, PowerShell lets you create modules, .NET classes (types), and Component Object Model (COM) objects. A discussion of these object types is beyond the scope of this article, but you should know that PowerShell can be extended far beyond what I show you here.

To create an object based on a .NET class, you need to specify the New-Object cmdlet, followed by the -TypeName parameter and the name of the class. In most cases, it's easiest to assign the new object to a variable, then use that variable to access the object's members. The following example creates an object based on the .NET class System.Net.NetworkInformation.Ping and assigns the object to the \$ping variable:

```
$ping = New-Object -TypeName Net.NetworkInformation.Ping
```

Notice that you don't need to specify *System* when you create the object. PowerShell knows where to look for the class.

That's all there is to creating the object. You can then access the object's members by using the `$ping` variable. In this case, to view the object's members, you need to pipe the variable's contents to the `Get-Member` cmdlet:

```
$ping | Get-Member
```

As you can see in Figure 5, the results show the object's members, as returned by the `Get-Member` cmdlet.

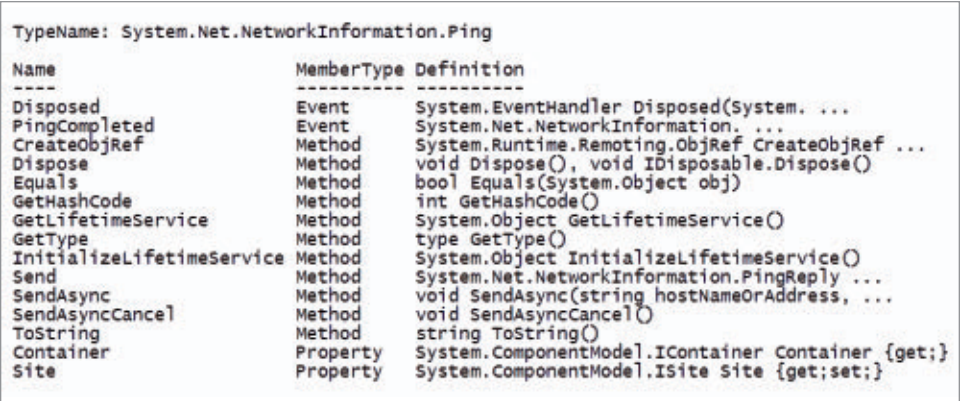


Figure 5
Creating an Object
Based on a .NET Class

As you would expect, the object is based on the `System.Net.NetworkInformation.Ping` class. Notice that the class supports a number of methods, one of which is `Send`. Let's take a closer look at that method by using the `Get-Member` cmdlet to return more detailed information:

```
$ping | Get-Member -Name Send | Format-List
```

As you can see, you use the `Get-Member` cmdlet with the `-Name` parameter followed by the `Send` method. You then pipe the information retrieved by `Get-Member` to the `Format-List` cmdlet, which returns

the results shown in Figure 6. Although the definition is quite extensive, it's essentially saying that the method lets you ping a specific website or IP address, as indicated by the part that reads *Send(string hostNameOrAddress)*.

Figure 6

Getting Information
About the Send
Methods

```

TypeName : System.Net.NetworkInformation.Ping
Name      : Send
MemberType : Method
Definition : System.Net.NetworkInformation.PingReply Send(string hostNameOrAddress),
System.Net.NetworkInformation.PingReply
Send(string hostNameOrAddress, int timeout),
System.Net.NetworkInformation.PingReply Send(ipaddress address),
System.Net.NetworkInformation.PingReply Send(ipaddress address,
int timeout),
System.Net.NetworkInformation.PingReply Send(string hostNameOrAddress,
int timeout, byte[] buffer),
System.Net.NetworkInformation.PingReply Send(ipaddress address, int
timeout, byte[] buffer),
System.Net.NetworkInformation.PingReply Send(string hostNameOrAddress,
int timeout, byte[] buffer),
System.Net.NetworkInformation.PingOptions options),
System.Net.NetworkInformation.PingReply Send(ipaddress address,
int timeout, byte[] buffer, System.Net.NetworkInformation.PingOptions
options)

```

The following example uses the Send method to ping *google.com*:

```
$ping.Send("google.com")
```

In this case, the Send method returns the results that you can see in Figure 7.

Figure 7

Using the Send
Method to Ping
google.com

```

Status      : Success
Address     : 173.194.33.37
RoundtripTime : 11
Options     : System.Net.NetworkInformation.PingOptions
Buffer      : {97, 98, 99, 100...}

```

Although this is a very simple example of what you can do with objects created from .NET classes, it demonstrates that you don't need to let cmdlet limitations prevent you from getting at the information you need. It helps to have knowledge about the various classes available in .NET, but you certainly don't need to be an expert in .NET development to take advantage of this powerful feature in PowerShell.

PowerShell's Object-Centric World

Objects are at the root of all PowerShell operations. The more you understand how objects work within the PowerShell environment, the better you'll be able to take advantage of PowerShell's flexibility and perform the variety of tasks that PowerShell supports. Not only can you use the PowerShell cmdlets to their fullest, but you can also create your own objects so you can push even deeper into your systems and network. And what you've seen here is only a glimpse of what you can do with objects in PowerShell. Objects give you inroads into the entire .NET Framework and well beyond. ■

CAN'T GET AWAY?

Get first-class education from your desk



Windows IT Pro offers FREE online events including webcasts, demos and virtual conferences. All events are brought to your computer live while being fully interactive.

Go to www.windowsitpro.com/events to see an up-to-date list of all online events.

Windows IT Pro

First-class education from the top experts in the industry.
Visit www.windowsitpro.com/events for a knowledge upgrade today!

Have a full plate on the live date? Don't sweat it! All online events are recorded and available 24/7.

Windows Server 2012: Building a Two-Node Failover Cluster

Learn how to add the Failover Clustering feature and use Failover Cluster Manager



**Michael
Otey**

is senior technical director
for *Windows IT Pro*
and *SQL Server Pro*.

Email



Windows Server 2012 has so many new features that it's tough to keep track of them all. Some of the most important new IT infrastructure building blocks are found in the improvements for failover clustering. Failover clustering originated as a technology that was designed to protect mission-critical applications such as Microsoft SQL Server and Microsoft Exchange, but since that time failover clustering has evolved into a high availability platform for several Windows services and applications. Failover clustering is part of the foundation for Dynamic Datacenter and technologies such as live migration. With Server 2012 and the improvements in the new Server Message Block (SMB) 3.0 protocol, failover clustering has been further expanded to enable continuously available file shares. For an overview of the all features in Server 2012 failover clustering, check out "[New Features of Windows Server 2012 Failover Clustering](#)."

I'll show you how to build a two-node Server 2012 failover cluster. First, I'll cover some the prerequisites and provide you with an overview of how the hardware environment, network, and storage are set up. Then, I'll dive into the details of how to add the Failover Clustering feature to Server 2012 and use Failover Cluster Manager to configure a two-node cluster.

Understanding the Failover Clustering Prerequisites

To build a two-node Server 2012 failover cluster, you need two systems running either the Datacenter or Standard edition of Server

2012. They can be physical systems or virtual machines (VMs). You can create clusters with VM nodes using either Microsoft Hyper-V or VMware vSphere. I'll be creating the cluster using two physical servers, but the cluster configuration steps are same regardless of whether the cluster nodes are physical or virtual. However, a key point is that the nodes should be similarly configured to enable the backup node to handle the workloads that might need to be supported in the event of a failover or live migration. You can see the overview of the components I used for my Server 2012 failover cluster in Figure 1.

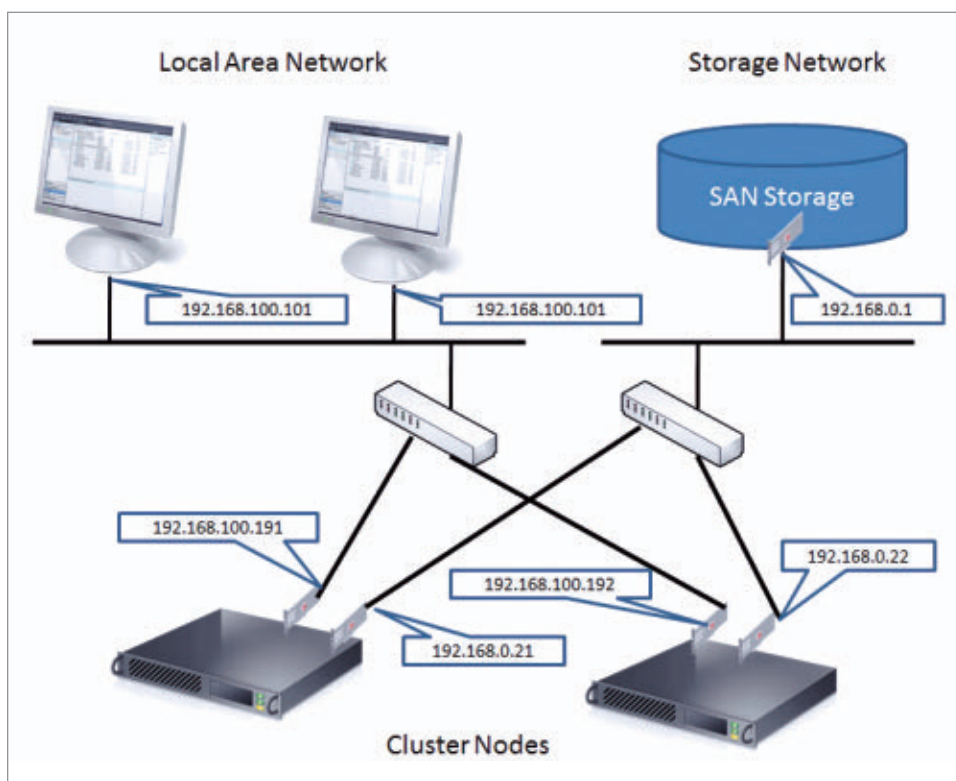


Figure 1

Reviewing the Cluster Components

A Server 2012 failover cluster requires shared storage, which can be an iSCSI, Serially Attached SCSI, or Fibre Channel SAN. In this example, I'm using an iSCSI SAN. When using this type of storage, you need to be aware of the following:

- Each server must be equipped with at least three NICs: one NIC dedicated to iSCSI storage connectivity, one NIC dedicated for cluster node communication, and one NIC for external network connections. If you're planning to use the cluster for live migration, you should consider having a fourth NIC dedicated to it. However, live migration can also occur over the external network connection—it'll just be slower. If you're using your servers for Hyper-V virtualization and server consolidation, you'll definitely want additional NICs to handle the VMs' network traffic.
- Faster is always better with networking, so the iSCSI connection should be running at a minimum of 1GHz.
- The iSCSI target must support the iSCSI-3 specifications, which include the ability to create persistent reservations. This is required by live migration. The iSCSI 3 standard is supported by almost all hardware storage vendors. If you're trying to implement a cluster in an inexpensive lab environment, you should make sure the iSCSI target software you're using supports iSCSI 3 and persistent reservations. Older versions of Openfiler didn't support this standard, but the new version of Openfiler with the [Advanced iSCSI Target Plugin](#) does support it. In addition, StarWind Software's [StarWind iSCSI SAN Free Edition](#) is fully compatible with Hyper-V and live migration. Certain versions of Windows Server can also act as an iSCSI target that's compatible with the iSCSI 3 standards. Server 2012 includes an iSCSI target. Windows Storage Server 2008 R2 includes support for iSCSI target software. Plus, you can download [Microsoft iSCSI Software Target 3.3](#), which runs on Windows Server 2008 R2.

You can find more details about how I configured the iSCSI storage for my failover cluster in the sidebar [“An Example of How to Configure iSCSI Storage.”](#) For more information about the requirements for failover clustering, you can check out [“Failover Clustering Hardware Requirements and Storage Options.”](#)

An Example of How to Configure iSCSI Storage

A **Windows Server 2012** failover cluster requires shared storage, which can be an iSCSI, Serially Attached SCSI, or Fibre Channel SAN. In my failover cluster, I configured an iSCSI SAN.

I started by creating three LUNs on the iSCSI SAN. I created one LUN for the cluster quorum and sized it at 520MB. I created another LUN for 10 virtual machines (VMs) and sized it at 375GB. I created the third LUN for a small test VM. I formatted all the LUNs using NTFS.

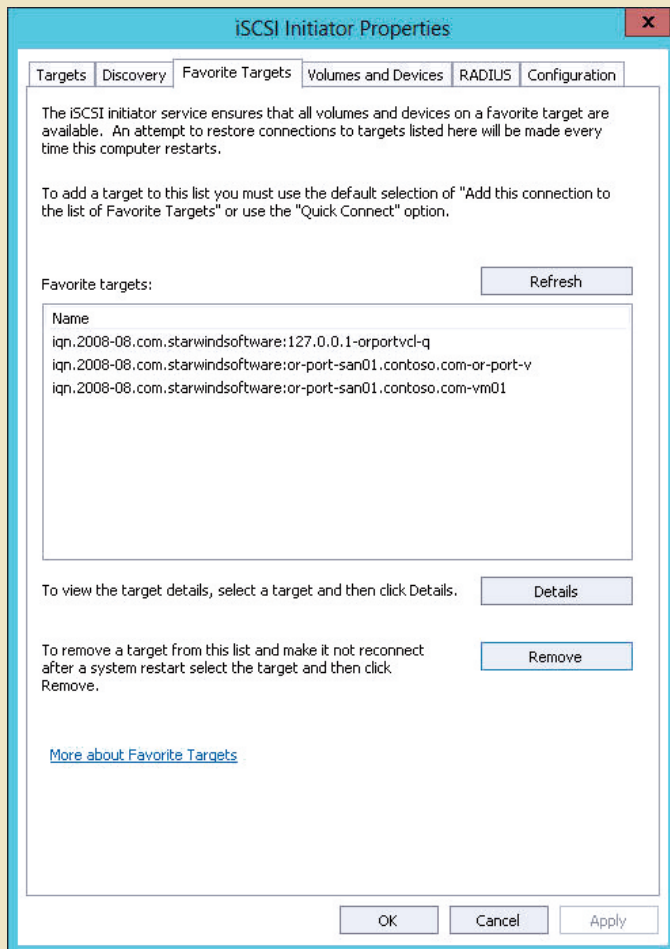


Figure A: Configuring the iSCSI Initiator

After creating the LUNs, I configured the iSCSI Initiator on both Server 2012 nodes. To add the iSCSI targets, I chose the iSCSI Initiator option on the Tools menu in Server Manager. On the Discovery tab, I clicked the Discover Portal button. This displayed the Discover Portal dialog box, where I entered the SAN's IP address (192.168.0.1) and iSCSI port (3260).

Next, I selected the Targets tab and clicked the Connect button. In the Connect To Target dialog box, I supplied the target name of the iSCSI SAN. I obtained this name from the SAN's properties. The name will vary depending on the SAN vendor, the domain name, and the names of the LUNs created. Besides supplying the target name, I selected the *Add this connection to the list of Favorite Targets* option.

After completing the iSCSI configuration, the iSCSI Initiator Targets tab was populated with the LUNs. To ensure that these LUNs would be automatically connected when Server 2012 starts, I made sure they were listed in the Favorite Targets tab, as shown in Figure A.

Finally, I assigned drive letters to the LUNs using the Microsoft Management Console (MMC) Disk Management snap-in. I choose Q for the quorum and W for the drive to be used for the VMs and Cluster Shared Volumes (CSVs). When assigning drive letters,

An Example of How to Configure iSCSI Storage (*continued*)

you first need to make the assignments on one node. Then, you need to bring the disks offline and make identical assignments on the second node. You can see the completed disk assignments for one of the nodes in Figure B. When you create the cluster, the drives will be shown as available storage.

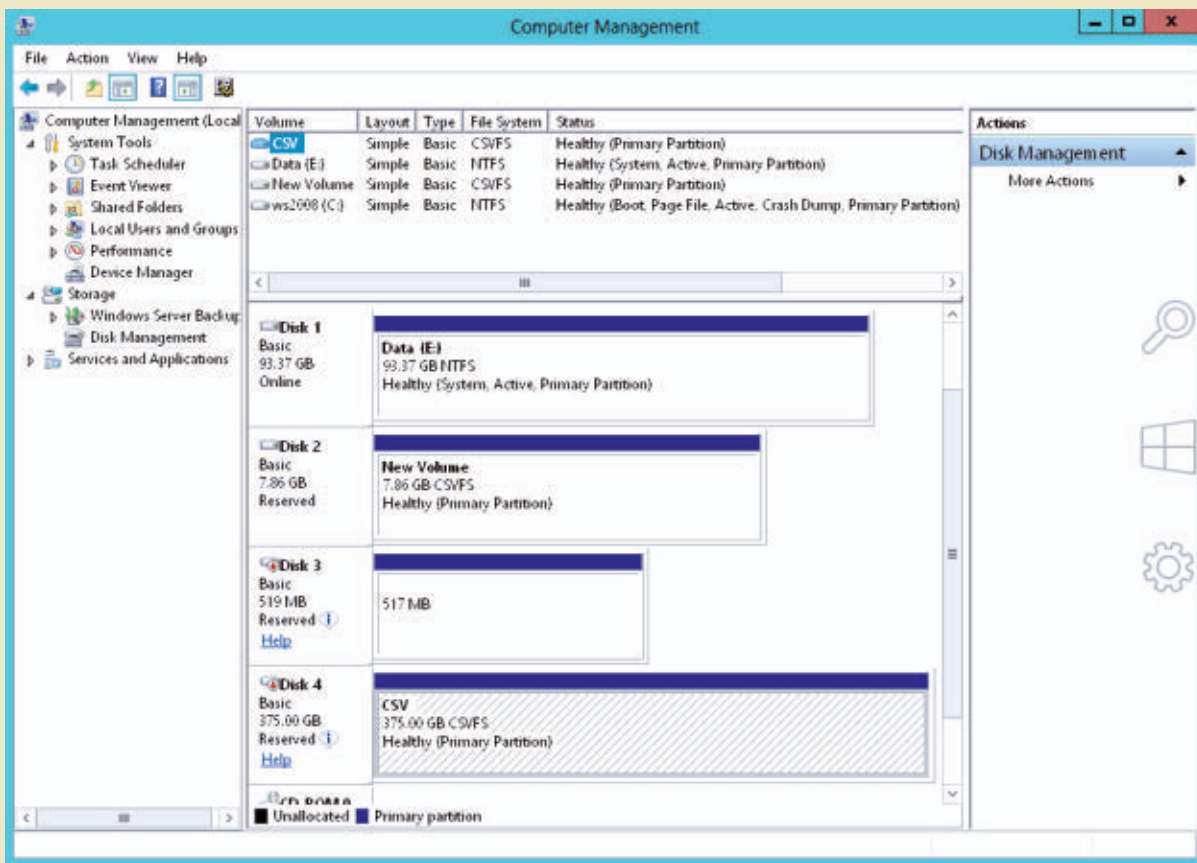


Figure B: Reviewing the Disk Assignments of a Node's iSCSI Drives

Adding the Failover Clustering Feature

The first step in creating a two-node Server 2012 failover cluster is to add the Failover Clustering feature using Server Manager. Server Manager automatically opens when you log on to Server 2012. To add the Failover Clustering feature, select Local Server and scroll down to

the ROLES AND FEATURES section. From the TASKS drop-down list, select *Add Roles and Features*, as shown in Figure 2. This will start the *Add Roles and Features* wizard.

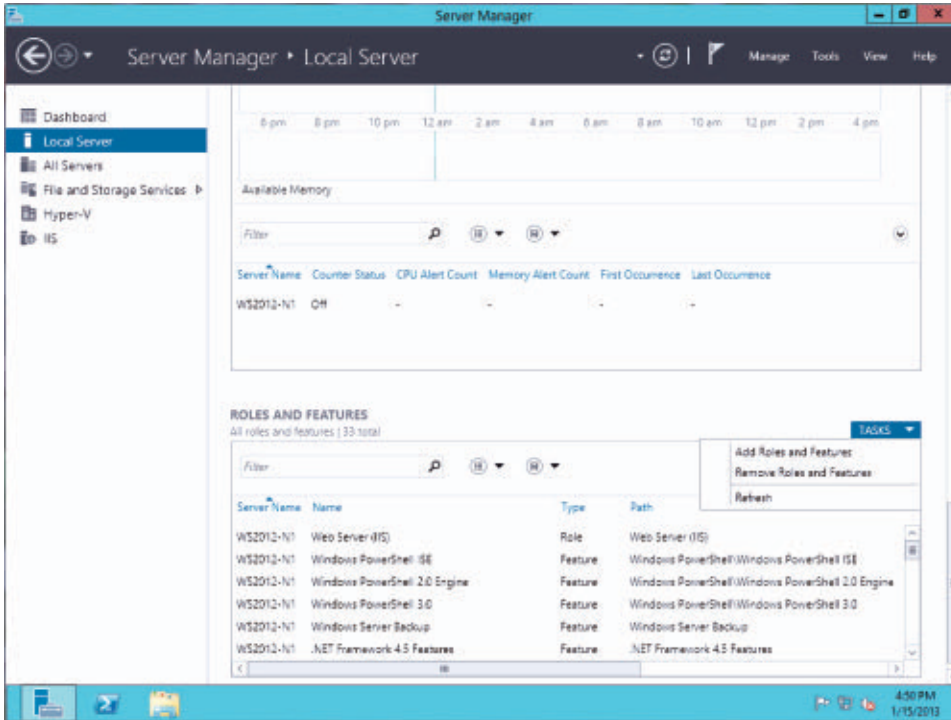


Figure 2

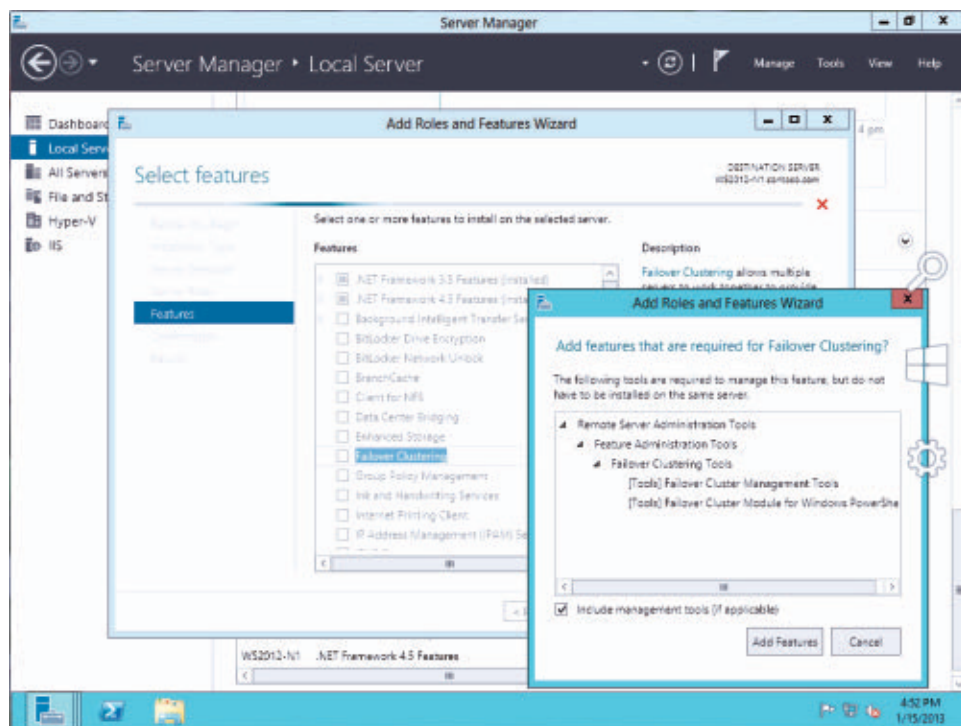
Starting the *Add Roles and Features* Wizard

The wizard opens with the *Before you begin* welcome page. Click Next to go to the *Select installation type* page, which basically asks if you're installing a feature on the local computer or installing a feature to a Remote Desktop service. For this example, select the *Role-based or feature-based installation* option and click Next.

On the *Select destination server* page, select the server on which you want to install the Failover Clustering feature. In my case, it was a local server named WS2012-N1. After selecting your local server, click Next to go to the *Select server roles* page. For this example, you won't be installing a server role, so click Next. Alternatively, you can click the Features link in the left menu.

On the *Select features* page, scroll through the Features list until you see Failover Clustering. When you click the box in front of Failover Clustering, the wizard displays a dialog box listing all the components that will be installed as part of this feature. As you can see in Figure 3, the wizard will install the Failover Cluster Management Tools and the Failover Cluster Module for Windows PowerShell by default. Click the Add Features button to return to the *Select features* page. Click Next.

Figure 3
Adding the Failover
Clustering Feature
and Tools



The *Confirm installation selections* page will list the Failover Clustering feature along with the management tools and PowerShell module. This page gives you a chance to go back and make any changes if needed. Clicking the Install button will begin the actual feature installation. After the installation completes, the wizard will end and Failover Clustering will be displayed in the ROLES AND FEATURES section of Server Manager. This process must be completed on both nodes.

Validating the Failover Clustering

After adding the Failover Clustering feature, the next step is to validate the configuration of the environment in which you'll create your cluster. To do this, you can use the *Validate a Configuration* wizard in Failover Cluster Manager. This wizard checks the hardware and software configuration of all the cluster nodes and reports on any issues that might prevent the cluster from being created.

To open Failover Cluster Manager, select the Failover Cluster Manager option on the Tools menu in Server Manager. In the Management pane, click the *Validate Configuration* link shown in Figure 4 to run the *Validate a Configuration* wizard.

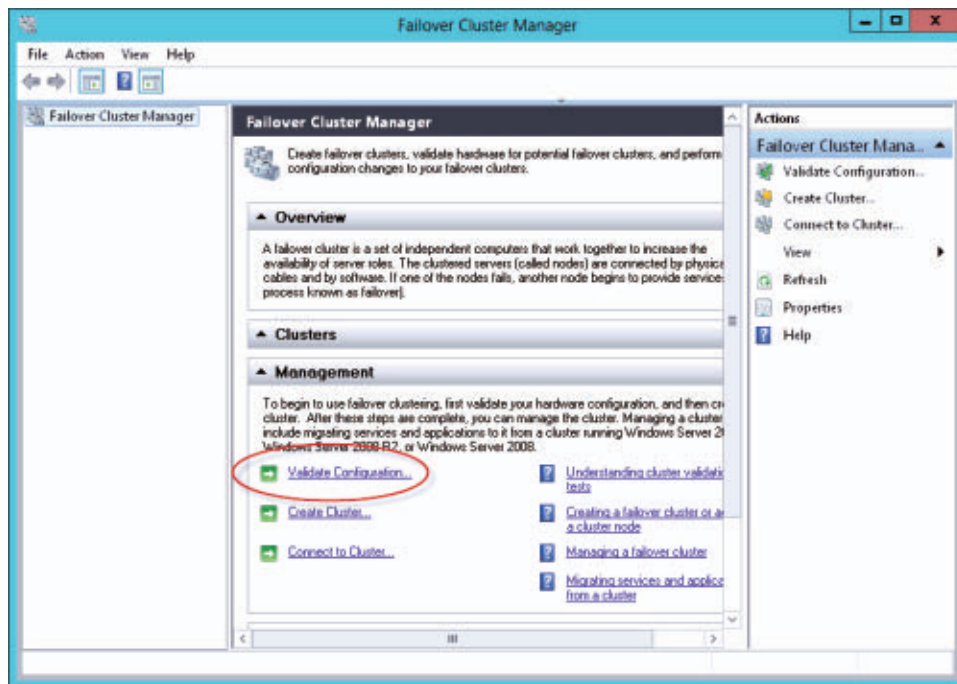


Figure 4

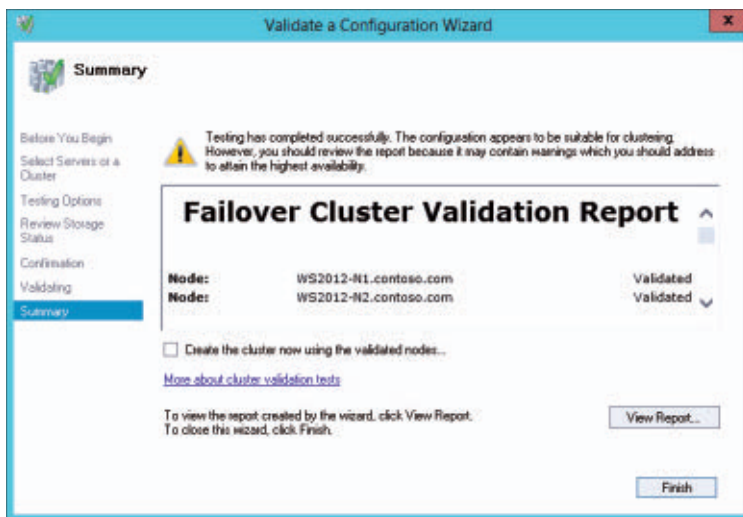
Starting the *Validate a Configuration* Wizard

The wizard first displays a welcome page. Click next to go to the *Select Servers or a Cluster* page. On this page, enter the names of the cluster nodes that you want to validate. I entered WS2012-N1 and WS2012-N2. Click Next to display the Testing Options page, where

you can select the tests that you want to run. You have the option to select specific sets of tests or to run all the tests. For at least the first time, I recommend that you select the option to run all the tests. Click Next to go to the Confirmation page, which shows the tests that will be run. Click Next to start the cluster validation testing process. The tests will check the OS level, network configuration, and storage of all the cluster nodes. A summary of the results are displayed when the test is finished.

If the validation tests succeed, you can create the cluster. Figure 5 shows the Summary screen for a successfully validated cluster. If errors are encountered during the validation tests, the validation report will display a yellow triangle for warning errors and a red X for severe errors. Warning errors should be reviewed, but they can be ignored. Severe errors must be corrected before the cluster can be created.

Figure 5
Reviewing the
Validation Report



Creating the Failover Cluster

At this point, you can create the cluster on any of the cluster nodes. I created the cluster on the first node (WS2012-N1).

To create a new cluster, select the Create Cluster link in either the Management pane or Actions pane, as Figure 6 shows.

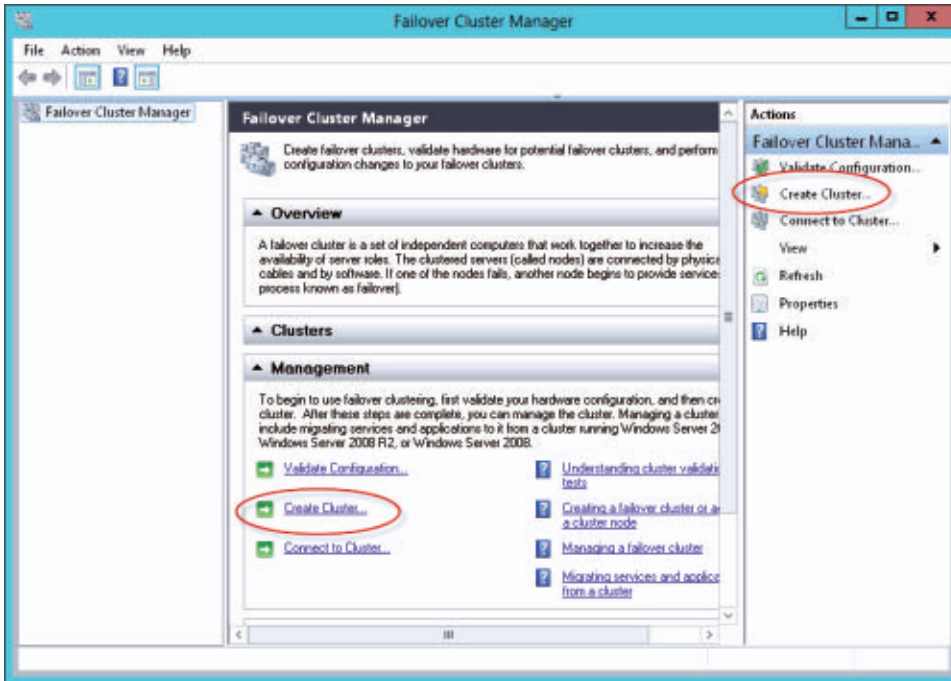


Figure 6
Starting the Create Cluster Wizard

This will start the Create Cluster wizard, which begins with a welcome page. Click Next to go to the Select Servers page shown in Figure 7. On this page, enter the names of all the cluster nodes, then click Next.

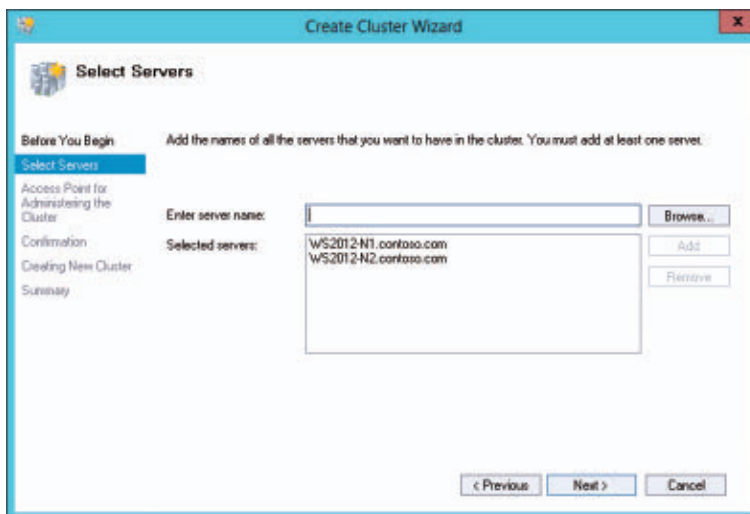
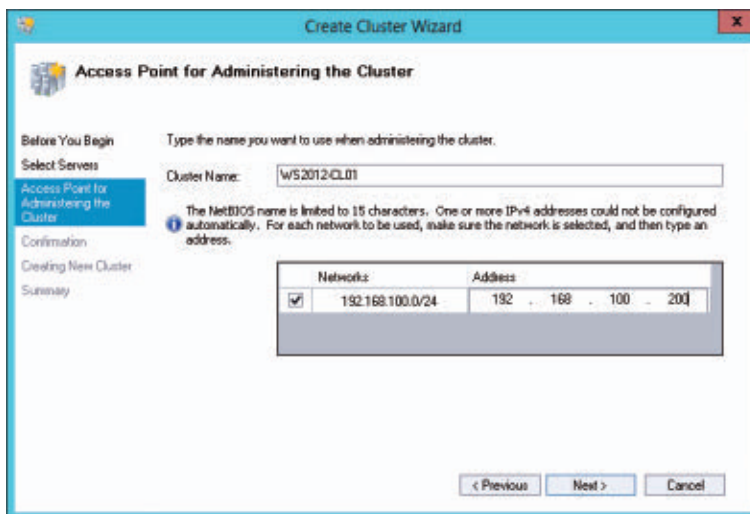


Figure 7
Selecting the Servers for the Cluster

On the *Access Point for Administering the Cluster* page, you specify your cluster's name and IP address, both of which must be unique in the network. In Figure 8, you can see that I named my cluster WS2012-CL01 and gave it an IP address of 192.168.100.200. With Server 2012, you can have the IP address of the cluster assigned by DHCP, but I prefer to use a statically assigned IP address for my server systems.

Figure 8
Configuring the
Cluster Access Point



After you enter the name and IP address, click **Next** to display the *Confirmation* page shown in Figure 9. This page lets you verify your cluster creation choices. If needed, you can page back and make changes.

Clicking **Next** on the *Confirmation* page creates the cluster on all of the selected clustered nodes. A progress page is displayed as the *Create Cluster* wizard goes through the steps of creating a new cluster. When it finishes, the wizard will display a *Summary* page that shows the configuration of the new cluster.

Although the *Create Cluster* wizard will automatically select the storage for your quorum, it often doesn't choose the quorum drive that you want. To check which disk is being used by the quorum, open the *Failover Cluster Manager* and expand the cluster. Then



Figure 9
Confirming the Cluster
Creation Selections

expand the Storage node and click the Disks node. The disks available to the cluster will be displayed in the Disks pane. The disk that the wizard selected for the cluster quorum will be listed under Disk Witness in Quorum.

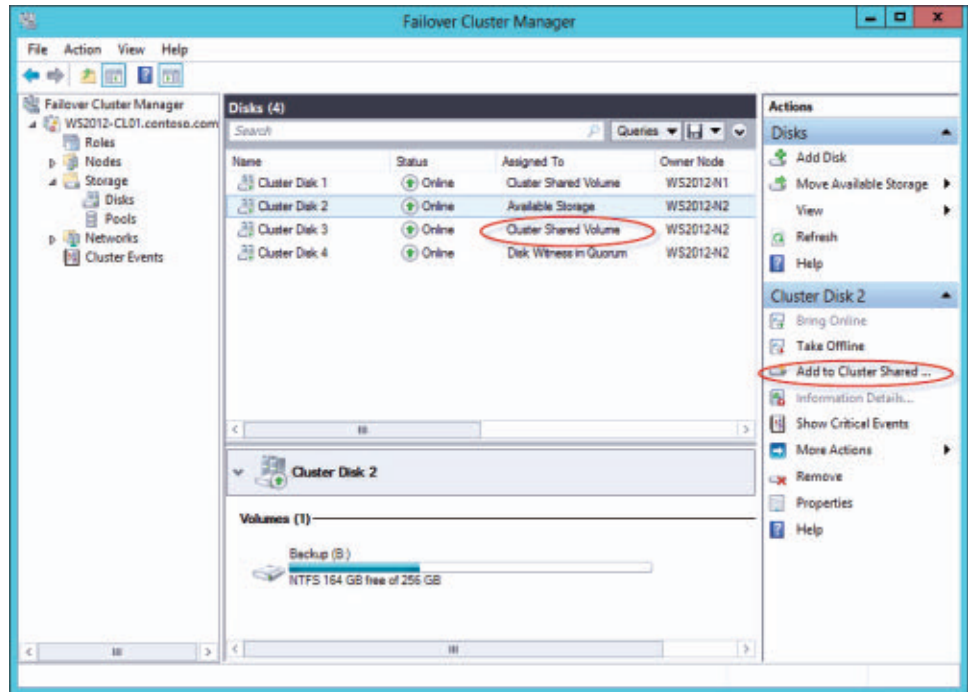
In my example, I used Cluster Disk 4 for the quorum. It was sized at 520MB, which is slightly larger than the quorum minimum of 512MB. If you want to use a different disk as the cluster quorum, you can change the quorum configuration by right-clicking the name of the cluster in Failover Cluster Manager, selecting More Actions, and choosing Configure Cluster Quorum Settings. This will display the Select Quorum Configuration wizard, which will let you change the cluster quorum.

Configuring Cluster Shared Volumes and the VM Role

Both nodes in my cluster have the Hyper-V role installed because I want to use the cluster for high-availability VMs supporting live migration. To help with live migration, the next step is to configure Cluster Shared Volumes (CSVs). Unlike Server 2008 R2 CSVs, Server 2012 CSVs are enabled by default. However, you still need to tell the cluster which storage should be used for the CSVs. To enable a CSV

on an available disk, expand the Storage node and select the Disks node. Next, select the cluster disk that you want to use as a CSV and click the *Add to Cluster Shared Volumes* link in the Failover Cluster Manager's Actions pane, as you see in Figure 10. That cluster disk's Assigned To field will then change from Available Storage to Cluster Shared Volume, as shown in Figure 10.

Figure 10
Adding a CSV



Behind the scenes, Failover Cluster Manager configures the cluster disk's storage for CSV, which includes adding a mount point in the system drive. In my example, I enabled CSVs on both Cluster Disk 1 and Cluster Disk 3, which added the following mount points:

- C:\ClusterStorage\Volume1
- C:\ClusterStorage\Volume2

At this point, the two-node Server 2012 cluster has been built and CSVs have been enabled. Next, you can install clustered applications

or add roles to the cluster. In my case, I'm building the cluster for virtualization support, so my next step is to add the Virtual Machine role to the cluster.

To add a new role, select the cluster name in Failover Cluster Manager's navigation pane and click the Configure Roles link in the Actions pane to launch the High Availability wizard. Click Next on the welcome page to go to the Select Role page. Scroll through the list of roles until you see the Virtual Machine role, as you see in Figure 11. Select that role and click Next.

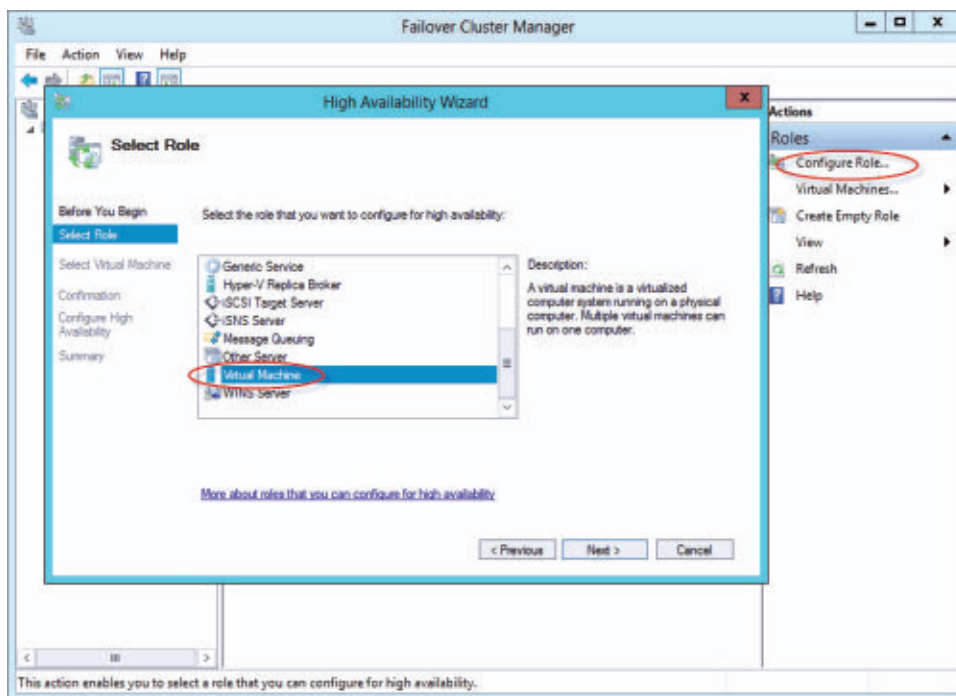
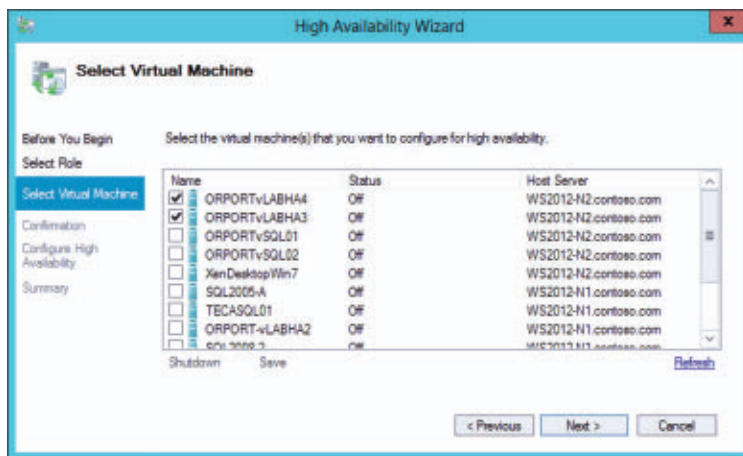


Figure 11
Adding a Virtual
Machine Role

On the Select Virtual Machine page that opens, all the VMs on all the cluster nodes will be listed, as shown in Figure 12. Scroll through the list of VMs and select the VMs that you want to be highly available. Then, click Next. After confirming your selections, click Next again to add the Virtual Machine roles to Failover Cluster Manager.

Figure 12
Selecting the VMs that
You Want to Make
Highly Available



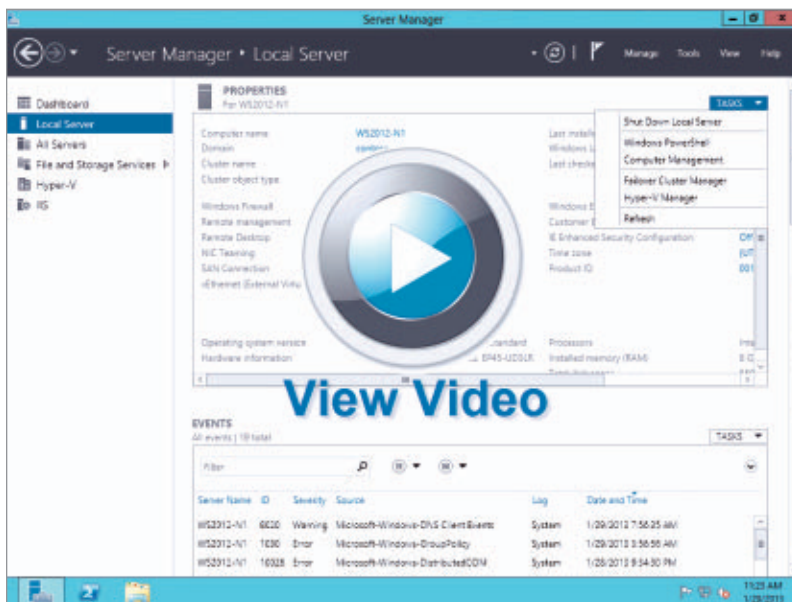
See This Process in Action

In this article, I explained how to create and configure a two-node Server 2012 cluster and how to add CSVs to the cluster and make

Video



Michael Otey discusses
how to configure a
two-node Windows
Server 2012 cluster



a VM highly available. To see the process in action, check out the accompanying video. ■



JOHAN ARWIDMARK



SEAN DEUBY



CHANDER DHALL



MARY JO FOLEY



DAN HOLME



MARK MINASI



MICHAEL OTEY



MARK RUSSINOVICH



ROD TRENT



PAUL THURROTT



For more than 13 years, IT/Dev Connections has been the premier training event for developers and IT professionals. IT/Dev Connections provides in-depth training on the technology platforms you're currently using, real-world solutions that will give you the competitive edge, and expert insight into how to plan for and implement the latest technologies. With more than 175 sessions to choose from, the conference offers training on HTML5, ASP.NET, Exchange, SQL Server, Windows OS, Windows Server, SharePoint, Visual Studio, Office 365, business intelligence, cloud, and all types of development.

BROUGHT TO YOU BY ▼



Go to **devconnections.com** to register

▼ Schedule at a Glance *(subject to change)*

MONDAY

7:30am	Registration Opens
9:00am – 4:00pm	Pre-conference workshops

TUESDAY

7:00am – 5:00pm	Registration Open
7:30am – 8:30am	Breakfast
8:30am – 9:30am	Keynote
9:30am – 10:00am	Break
10:00am – 11:15am	Sessions
11:15am – 11:45am	Break
11:45am – 1:00pm	Sessions
1:00pm – 2:30pm	Lunch
2:30pm – 3:45pm	Sessions
3:45pm – 4:15pm	Break
4:15pm – 5:30pm	Sessions
5:30pm – 7:30pm	Welcome Reception

WEDNESDAY

7:00am – 5:00pm	Registration Open
7:30am – 8:30am	Breakfast
8:30am – 9:30am	Keynote
9:30am – 10:00am	Break
10:00am – 11:15am	Sessions
11:15am – 11:45am	Break
11:45am – 1:00pm	Sessions
1:00pm – 2:30pm	Lunch
2:30pm – 3:45pm	Sessions
3:45pm – 4:15pm	Break
4:15pm – 5:30pm	Sessions

THURSDAY

7:00am – 5:00pm	Registration Open
7:30am – 8:30am	Breakfast
8:30am – 9:30am	Keynote
9:30am – 10:00am	Break
10:00am – 11:15am	Sessions
11:15am – 11:45am	Break
11:45am – 1:00pm	Sessions
1:00pm – 2:30pm	Lunch
2:30pm – 3:45pm	Sessions
3:45pm – 4:15pm	Break
4:15pm – 5:30pm	Sessions

FRIDAY

7:30am	Registration Opens
9:00am – 4:00pm	Post-conference workshops

Make Connections the conference you bring your whole team to this year!

NEW LOOK!
same great conference
you know and love!

For more than 13 years, Connections has been the go-to Microsoft training conference for IT Professionals and Developers.

IT/Dev Connections is the only in-person training conference you need to attend this year. Join thousands of Microsoft developers, DBAs and IT Professionals and stay on top of the latest in:

- ASP.NET
- HTML5
- Visual Studio
- SQL Server
- Windows Server
- SharePoint
- Exchange
- and much, much more.

Register with the
**All Access
VIP PASS
and Save!**



CONFERENCE AT A GLANCE

3 full days of educational sessions, with pre- and post-conference workshops

Train with 100+ Microsoft & industry experts

More than 175 in-depth technical, how-to sessions

Exciting and entertaining networking opportunities

Discover new products and services in the partner expo hall

Ask YOUR pressing technical strategy questions, hear diverse opinions from our expert speakers and develop your company's plan for the future!

HIGHLY ANTICIPATED SPEAKERS



MARY JO FOLEY
Microsoft Expert
AllAboutMicrosoft.com



MIGUEL DE ICAZA
CTO, Co-founder of the
GNOME and MONO projects
Xamarin



MARK MINASI
Senior Contributing Editor
Windows IT Pro



MARK RUSSINOVICH
Senior Contributing Editor
Windows IT Pro



PAUL THURROTT
Senior Technical Analyst
Windows IT Pro

MORE EXPERT SPEAKERS

This year's conference speakers were chosen out of more than 500 stellar session entries, guaranteeing the best-of-the-best are speaking at Connections 2013.



JOHAN ARWIDMARK
Knowledge Factory



ITZIK BEN-GAN
SQL Server Pro



JES SCHULTZ BORLAND
Brent Ozar Unlimited



ANDREW CONNELL
Critical Path Training



BRIAN DESMOND
Moran Technologies



SEAN DEUBY
Windows IT Pro



CHANDER DHALL
Microsoft MVP



TIM FORD
Spectrum Health



SCOT HILLIER
SharePoint MVP



DAN HOLME
Microsoft MVP



STEVE JONES
SQL Server Central



MICHAEL OTEY
SQL Server Pro



BRENT OZAR
Brent Ozar Unlimited



JEREMIAH PESCHKA
Microsoft MVP



TONY REDMOND
Windows IT Pro



PAUL ROBICHAUX
Windows IT Pro



LORYAN STRANT
Office 365 MVP



JEREMY THAKE
Microsoft MVP



ORIN THOMAS
Windows IT Pro



ROD TRENT
Windows IT Pro



RANDY WILLIAMS
AvePoint



NATHAN WINTERS
Microsoft

...and many more!

See website for updated speakers and sessions.

(Speakers are subject to change.)

175+ sessions

*Pick and choose
from more than 175+
sessions to optimize
your training
experience.*

HTML5 | Visual Studio | ASP.NET
Windows | Exchange | Cloud
SQL Server | SharePoint

Sessions subject to change.

*Go to www.devconnections.com to see the most
up-to-date schedule.*

SQL Server

- How Active Directory Affects SQL Server
- Manage Your Shop with CMS and Policy Based Management
- Inside the Query Optimizer
- Data Internals Deep Dive
- T-SQL Querying and Query Tuning Enhancements in the Latest Major Releases of SQL Server
- Efficient Interval Management in SQL Server
- Practical Uses of Window Functions
- Use Dynamic Management Views to Diagnose SQL Server Performance Issues
- Hardware 201: Selecting and Sizing Database Hardware for OLTP Performance
- SQL Server Transaction Log Internals
- SQL Server Index Internals
- SQL Server 2012 in a Highly Available World
- Table Indexing for the .NET Developer
- SQL Server Table Partitioning from the Ground Up
- Using Power View and Hadoop to Unlock Hidden Markets
- Build Your Own SQL Server Private Cloud
- Surviving Your Peak Database Load
- Troubleshooting SQL Server with SysInternals Tools

- Code-Less Securing of SQL Server
- Shortcuts to Productivity in SQL Server Management Studio
- From Zero to Hero: A Case Study in Reducing Extremely High I/O on a SQL Server System
- SQL Server Optimization: Tuning the Hardware Subsystems
- What DBAs Need to Know About Hekaton
- Using FullText Search with Office Documents and PDFs
- Encryption in SQL Server
- Collecting and Analyzing File and Wait Statistics
- Maximizing Plan Reuse
- Practical Performance Monitoring in SQL Server
- Windows Azure SQL Database Troubleshooting and Query Tuning
- Windows Azure SQL Database for the DBA
- Make Your Queries Fly With Columnstore Indexes
- Using BIML as an SSIS Design Patterns Engine
- Hacking the SSIS 2012 Catalog
- SSIS Design Patterns
- Reduce, Reuse, Recycle: Automating Your BI Framework
- BI Security Best Practices
- From Reporting Services Rookie to Rock Star
- Maximizing SSIS Package Performance
- Cleaning Up Dirty Data in SSIS
- Choosing Between SSAS 2012 Multidimensional and Tabular
- PowerPivot to SSAS 2012 Tabular
- Managing SQL Server Performance with Extended Events
- Leveraging the Plan Cache for Performance Tuning
- Improve the Performance of Your T-SQL Queries by Changing Your Habits
- Creating SSRS Reports Efficiently Through Best Practices
- Manage SQL Server Efficiently with PowerShell Remoting
- Manage SQL Server 2012 on Windows Server Core with PowerShell
- TempDB Performance Troubleshooting and Optimizing
- Page Latches for Mere Mortals

- Creating Highly Performant, Scalable Websites: From SPA to Backend
- Architecting Device-driven Mobile Web Solutions
- Building Secured, Scalable, Low-Latency Web Applications with the Windows Azure Platform
- Writing Next Generation JavaScript with TypeScript
- Creating Data-Driven Mobile Web Apps with ASP.NET MVC and jQuery Mobile
- Touch-Enabled and Data Connected Sites in Hours, Not Weeks
- Unit Testing ASP.NET MVC
- Software Gardening
- Branches and Merges Are Bears, Oh My!
- Exploring Domain-Driven Design Implementation Patterns in .NET
- Behavior-Driven Development: Turning User Stories into Executable Specifications
- CQRS: Crack for Architecture Addicts?
- Get More Bang for Your Windows Azure Buck!
- Grokking Caching
- From Developer to Architect: 10 Things You Must Know
- Everything You Need to Know About Trends in Application Development
- Linq to Everything
- Asynchronous Programming with Async and Await
- Build your first Angular Web Application
- Building Games for Windows 8 – Using GameMaker
- Best Practices for Building Windows Phone and Windows 8 Applications
- Domain-Driven Design, CQRS, and Event-Sourcing for the Busy Developer
- Do's and Don'ts of Software Projects
- Fast Facts of Social Network Programming
- Creating Data-Driven HTML5 Applications
- Building End-to-End Web Apps Using TypeScript
- jQuery Fundamentals
- Automating Windows Azure from the Command Line
- Caching in Azure: There's More to That Than Azure Caching
- Debugging the Web with Fiddler
- IIS for Developers
- Doing It Right: Continuous Delivery Doesn't Have to Suck

- Building a Windows 8 App from Scratch
- Git for Visual Studio Developers
- From Manual Testing to Automation with Visual Studio ALM
- New Features in Visual Studio 2013 and TFS 2013
- Strategies for Refactoring and Testing Legacy Code
- Connecting the Dots: Using HTML5, jQuery, and Web API Together
- Advanced Debugging with WinDbg and SOS
- Task and Data Parallelism: Real-World Examples
- .NET Garbage Collection Performance Tips
- How to (Remote) Control Office 365 with Windows Azure
- Using Async in Your Mobile Apps
- A .NET Developer's Guide to Mobile Apps
- Cloud Data for the Everyday Developer
- Building Solutions in the Cloud with Apps for Office
- Developing Professional Solutions for Office 2013 and Outlook
- Doing Modern Web in the Enterprise
- Web Performance Optimization for Modern Web Applications
- Debugging and Testing JavaScript in Today's Browsers
- Developing Neural Networks with Visual Studio
- Introduction to iPhone Programming with C#, .NET, and Xamarin.iOS
- Introduction to Android Programming with C#, .NET, and Xamarin.Android
- Introduction to Mobile Web with HTML5
- Essential Typescript
- Unit Testing Web Development
- Customizing the SharePoint 2013 user interface with JavaScript
- JavaScript for Windows 8 Developers, Part 1 and 2
- JavaScript Testing - An Introduction
- Building Cross-Platform Mobile Applications with PhoneGap, Part 1 and Part 2
- Developing with the SharePoint 2013 App Model
- Creating Line-of-Business Apps in HTML5 and MVC/Web API
- Everyday Bootstrap
- Simplify Your API: Creating Maintainable and Discoverable Code

- Introduction to PowerShell for the Anxious IT Pro
- SharePoint Performance: Best Practices from the Field
- Who Says You Can't Do Records Management in SharePoint?
- Top 10 New ECM Features in SharePoint 2013
- Dan Holme's SharePoint 2013 MasterClass: SharePoint Installation and Configuration, From Bare Metal to Farm
- Implementing End-to-End SharePoint Governance
- Best Practices for Role-Based Management of Users, Groups, Permissions, Service Accounts, and Administrative Delegation
- Developing Search Applications in SharePoint 2013
- Developers Approach to Social Applications with SharePoint 2013
- The Only Way to Go is Up! Upgrading to SharePoint 2013
- Optimizing and Accelerating Your SharePoint Farm
- 0 to 60: Apps for Office and SharePoint
- Migrating SharePoint Solutions to Apps for SharePoint
- What Options Do Non-Developers Have in SharePoint 2013?
- Create Powerful SharePoint Designer 2013 Workflows in Office 365 and On-Premises
- Office 365: Introduction to SharePoint Online Development
- Surfacing Your Azure External Data Using BCS in SharePoint 2013 with Alerts
- Data Visualization with SharePoint and SQL Server
- Extending the Business Process Management Features of Office 365

- Exchange Online: Real-World Migration Challenges and Solutions
- Notes from the Field: Running a 500,000-Mailbox On-Premises Exchange Server Deployment
- Exchange Server 2013 Site Resiliency
- Managed Availability: Ensuring the End User Experience
- Data Loss Prevention in the Real World
- Hybrid and SSO Deployment with the New Office 365 (Wave 15)
- The Tao Of Exchange Server 2013 Sizing
- Virtualizing Exchange Server 2013: Why Not?
- From Zero to Hero: PowerShell for Exchange Server Boot Camp
- Exchange Server 2013 Unified Messaging Deep Dive
- Better Together: Integrating Exchange Server 2013 and Lync Server 2013
- Migrate to Modern Public Folders the Worry-Free Way
- Troubleshooting Modern Public Folders: A DIY Guide
- Apples to Apples: Comparing Office 365 to the Competition
- How Does Microsoft Secure My Email with Office 365?
- CAS 2013 – Why It Is 3 Better Than CAS 2010 and 6 Better Than 2007
- Building a Hybrid Configuration with Exchange Server 2013 in (Less Than) 75 Minutes
- How-to: Load Balancing Exchange Server 2013
- Exchange ActiveSync: Taming the Beast
- Exchange Server 2013 Backup, Restore, and Recovery
- High Availability in Exchange: A Recipe for Success?

- Managing Third-Party Updates with System Center 2012 Configuration Manager SP1
- Migrating from Configuration Manager 2007 to Configuration Manager 2012
- Developing Hydration Kits – IT Pro Automation at Its Best!
- A Geek's Guide to USMT 5.0
- Configuration Manager 2012 SP1 OS Deployment
- Using Windows Azure Infrastructure as a Service as Your Data Center
- What's New in Windows Server 2012 Hyper-V
- The WHY of Configuration Manager
- Hierarchy Simplification with Configuration Manager 2012
- Deploying and Managing Virtual Applications and Settings with System Center Configuration Manager and MDOP
- Deploying and Managing Virtual Applications and Settings with Active Directory Domain Services and MDOP
- Smoothing the Kinks for a Seamless User Experience with Microsoft UE-V
- Deploying Your Office in the Cloud with Office 365
- Windows Server 2012 Advanced Troubleshooting Workshop
- Troubleshooting Group Policy in Windows Server 2012
- Managing Public Cloud Infrastructure with PowerShell
- Manage Server 2012 Like a Pro or, Better, Like an Evil Overlord!
- Windows "Next:" Will Blue Make You Blue?
- AppLocker: Your Solution for Application Smackdown!
- Deploy Office 2010 or Office 2013 Using Group Policy (It CAN Be Done!)
- Windows Installer Survival Guide for System Center Configuration Manager Admins
- System Center Configuration Manager Software Update Zen
- Using Orchestrator to Integrate with Azure IaaS
- Become an Orchestrator Master
- Hyper-V Best Practices
- Microsoft Windows PowerShell Remoting In-Depth
- State-Based Administration of the Modern Enterprise
- Configuration Manager for UNIX and Mac — Myths and Realities
- Windows Intune Overview
- Managing Devices in the Cloud with Windows Intune

Mandalay Bay Resort & Casino

Network with your colleagues at Mandalay Bay Resort & Casino! There's so much to do, you'll never have to leave this 4-star resort!

HOTEL ACCOMMODATIONS

Mandalay Bay Resort & Casino
3950 Las Vegas Blvd. South, Las Vegas, NV
SPACE IS LIMITED so reserve your room early.
Call: 877-632-9001 and reference IT/Dev Connections
Room Block Rates Expire September 15, 2013

ATTIRE

The recommended dress for the conference is casual and comfortable. Please bring along a sweater or jacket, as the ballrooms can get cool with the hotel's air conditioning.

TAX DEDUCTIONS

Your attendance to a DevConnections conference may be tax deductible. Visit www.irs.ustreas.gov. Look for topic 513 — Educational Expenses. You may be able to deduct the conference fee if you undertake to (1) maintain or improve your skills required in your present job; (2) fulfill an employment condition mandated by your employer to keep your salary, status, or job.

GROUP DISCOUNTS

Register individuals from one company at the same time and receive a group discount (10% off registration. Not to be combined with other discounts or offers).



Registration & Cancellation Policy: Registrations are not confirmed until payment is received. Cancellations before August 2, 2013, must be received in writing and will be refunded minus a \$100 processing fee. After August 2, 2013, cancellations and no shows are liable for full registration; it can be transferred to the next Conference within 12 months or to another person. You may transfer this registration to a colleague by notifying us before the start of the event. Please inform us if you have any special needs or dietary restrictions when you register. The Conference Producers reserve the right to cancel the conference by refunding the registration fee. Producers can substitute speakers and topics and cancel sessions without notice or obligation. Updates will be posted on our website at www.DevConnections.com.

Notes & Policies: Tape recording, video recording and photography are not allowed at any session. Conference producers will be taking candid pictures of events and reserve the right to reproduce. By attending this conference you agree to this policy. Microsoft, Microsoft .NET, ASP.NET, Visual Studio, Microsoft SQL Server, Exchange, SharePoint and Windows are either trademarks or registered trademarks of Microsoft Corporation. All other trademarks are property of their owners.

Conference Registration

Full Conference Registration Includes Keynote on October 1, 2013, through Closing Session October 3, 2013.

Name	Discount Code
Company	Title
Street Address (Required to ship materials)	
City/State/Postal Code	
Country	
Phone	Fax
E-Mail (required)	

Online:
devconnections.com

Email:
Info@devconnections.com

Phone:
888.899.0130

Fax:
800.766.5367

Mail:
Penton Media
DevConnections
24654 Network Place
Chicago, IL 60673

Check the conference track you are registering for.
NOTE: you can attend any of the co-located conference tracks for no additional charge.

- ☐ Dev Connections
- ☐ Windows Connections
- ☐ SQL Server Connections
- ☐ SharePoint Connections
- ☐ Exchange Connections
- ☐ ALL ACCESS VIP PASS..... \$2,695
- ☐ BASIC REGISTRATION \$1,695
- ☐ Pre-Conference Workshops
Monday, September 30, 2013..... \$499
- ☐ Post-Conference Workshops
Friday, October 4, 2013..... \$499

TOTAL

Payment Information:

☐ CHECK
*(payable to Penton Media) All payments must be in US currency.
Checks must be drawn on a US bank.*

☐ CREDIT CARD:
☐ VISA ☐ MASTERCARD ☐ AMEX

--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Credit Card No.

--	--	--	--	--	--

Expiration Date

Cardholder's Signature

Cardholder's Name (print)

Top 5 Features in System Center 2012 Configuration Manager SP1

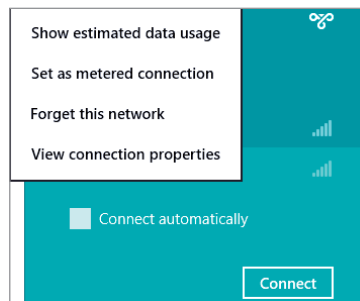
Find out why this service pack offers more than expected

A service pack is often a collection of hotfixes and maybe one or two new features. Microsoft System Center 2012 Configuration Manager Service Pack 1 (SP1) is an exception to the rule. It's full of so many new features that it should be called a *feature* pack. I'll summarize my personal top 5 favorites and explain why the SP1 release is a must-install.

1. Support for Windows 8 and Windows Server 2012

Configuration Manager now supports all site systems that run on [Windows Server 2012](#), including the primary site server. There is also full feature support for [Windows 8](#). And some new features—such as metered networks, user data and profiles, and modern style applications—are supported *only* on Windows 8.

Metered networks. Metered networks in Windows 8 protect users who connect via connections such as 3G or 4G (which have a cost associated with data transfer) from getting a huge bill from their cell providers. With Configuration Manager, you can control the download behavior for each deployment and each device, as Figure 1 shows.



Kent Agerlund

is a Microsoft System Center 2012 Configuration Manager MVP who works as senior System Center architect, trainer, event speaker, and author. For the past four years, he has been on the road with his Mastering System Center 2012 Configuration Manager class.

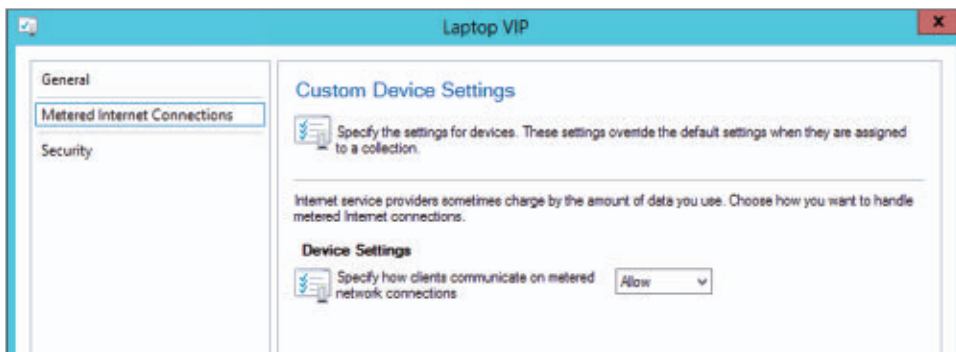


Email

Figure 1
Configuring a Metered Connection in Windows 8

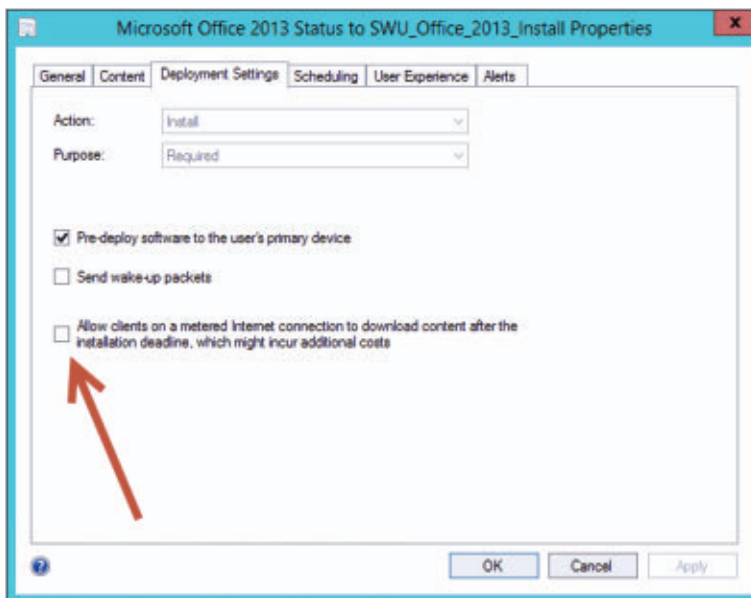
1. To control who can download while connected to a metered connection, create a new custom client device setting in the Administrator workspace. Choose Metered Internet Connections. Choose Allow under Device Settings, as Figure 2 shows.

Figure 2
Configuring
Support for Metered
Connections



2. Close the custom settings and deploy them to a collection of laptops.
3. For each deployment, you can configure whether to allow the download and installation on metered networks, as Figure 3 shows.

Figure 3
Controlling the
Deployment of
Metered Connections



User data and profiles. User data and profiles have been around in Microsoft environments for a long time. However, they were previously accessible only through Active Directory (AD) and Group Policy Objects (GPOs). Now, Configuration Manager allows administrators to manage and report on user profile settings such as folder redirections, offline files, and roaming profiles. The main benefits of controlling user data and profiles in Configuration Manager instead of in AD are the flexibility and the ability to configure the feature in a reporting-only mode. As with any other deployment, you simply create the settings and deploy them to a collection. Working with collections doesn't require users to log off or restart computers.

1. Start the user data and profiles process in the *Asset and Compliance* workspace.
2. Choose Compliance Settings, *User Data and Profiles*. Click Create User Data Profiles Configuration Item on the Ribbon.
3. Select the settings that you want to control (as Figure 4 shows) and click Next.

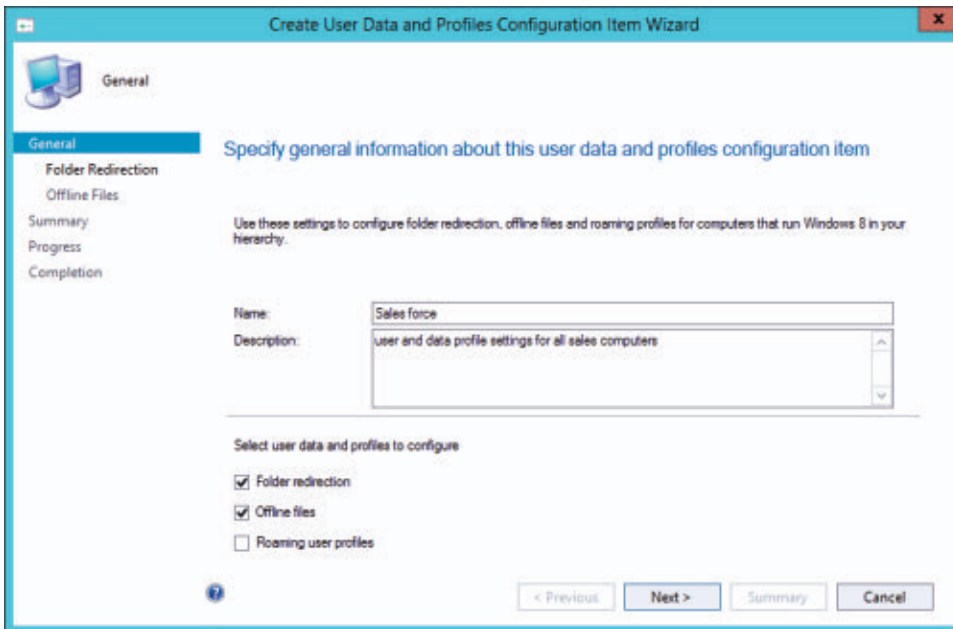
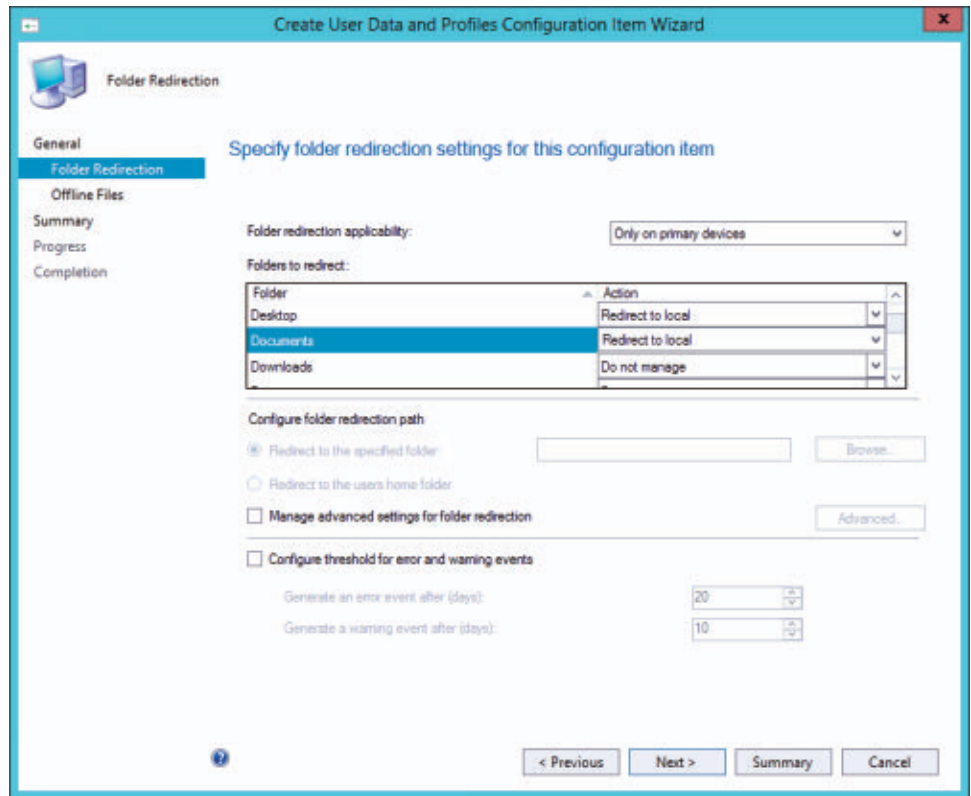


Figure 4
Monitoring and
Controlling User Data
and Profiles

Figure 5
Controlling Folder
Redirection Settings



4. On the Folder Redirection page, which Figure 5 shows, you can configure a few settings, control which device the settings will apply to, and configure thresholds for alerts and warnings.
5. The next setting to control is offline files, as Figure 6 shows. Again, most of these settings can also be controlled by Group Policy, but that method doesn't give you the monitoring option or flexibility of using collections.
6. Finish the wizard and then click Deploy on the Ribbon. In the *Deploy User Data and Profiles Configuration Item* dialog box, assign a user collection and then set automatic remediation and compliance thresholds. See Figure 7 for an example of deploying user data and profile settings.

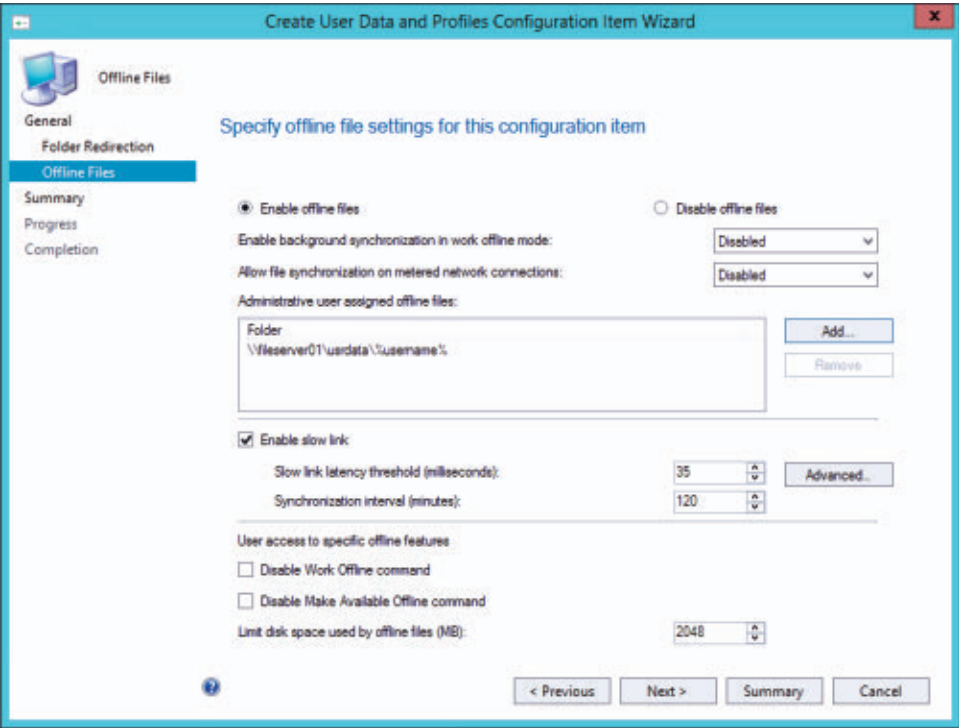


Figure 6
Configuring Offline
Settings

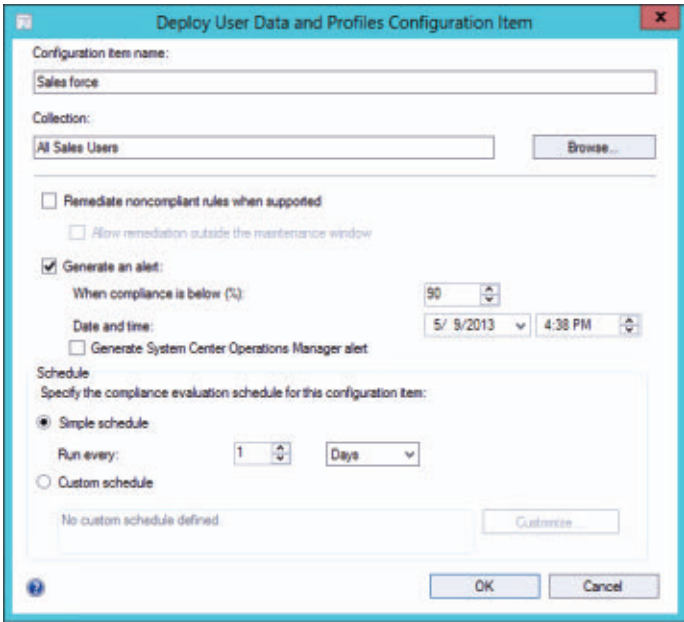


Figure 7
Deploying User Data
and Profile Settings

2. Windows PowerShell Support

Long have we been waiting for true PowerShell support in Configuration Manager. Finally, the wait is over. Microsoft envisions that all features in the Configuration Manager console will be available as PowerShell cmdlets. Microsoft will continue to add more PowerShell cmdlets in upcoming cumulative update releases. The most recent release, CU1, adds 40 new cmdlets, bringing the total number to 511.

You can launch PowerShell from within the Configuration Manager console. Doing so also launches the Configuration Manager module. Or you can launch PowerShell and manually launch the Configuration Manager PowerShell module.

1. To open the Configuration Manager Administrator console, choose Start, Connect via Windows PowerShell.
2. In the PowerShell console, enter A (to always trust the publisher) and press Enter.
3. Enter

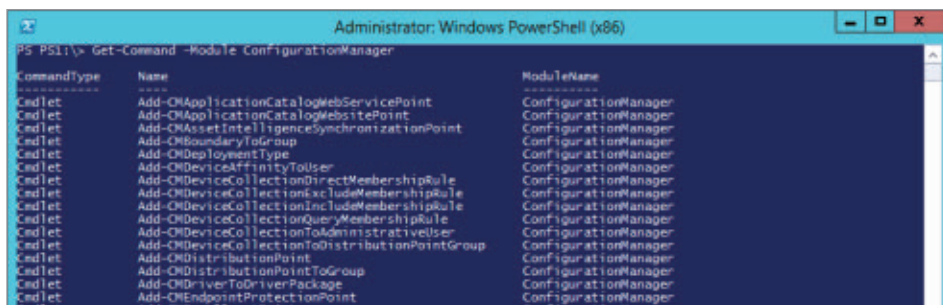
```
Get-Command -Module ConfigurationManager
```

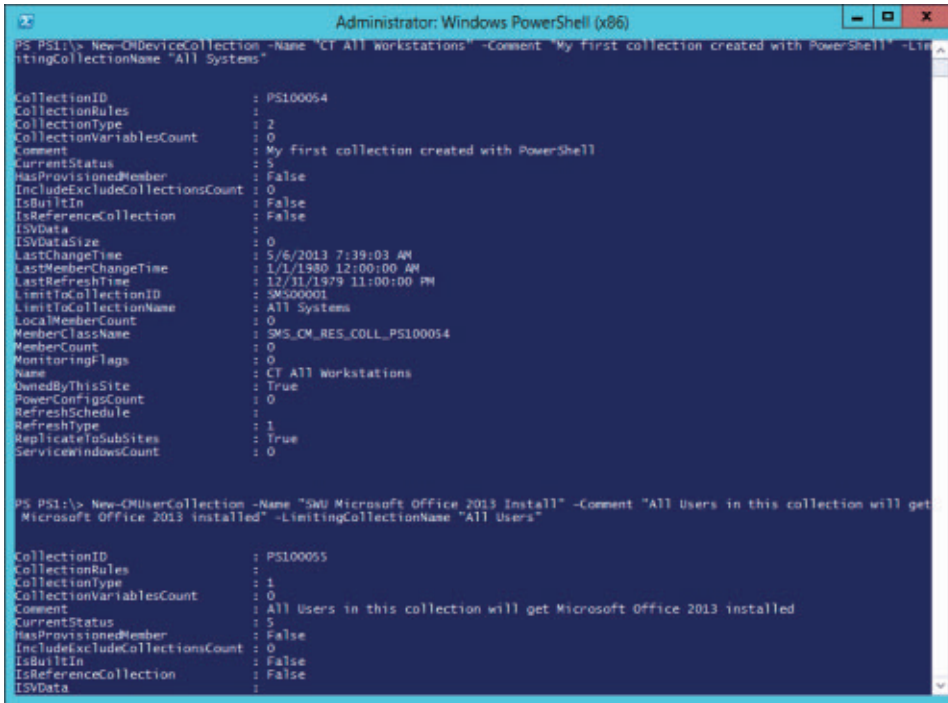
and press Enter to list all the Configuration Manager cmdlets, as Figure 8 shows.

4. The built-in cmdlets will make your life as a Configuration Manager administrator easier. Figure 9 shows how to create two new collections: a device collection and a user collection. Note that each collection is limited to another collection.

Figure 8

Listing Built-In Cmdlets





```

Administrator: Windows PowerShell (x86)

PS PS1:\> New-CMDeviceCollection -Name "CT All Workstations" -Comment "My first collection created with PowerShell" -LimitingCollectionName "All Systems"

CollectionID                : PS100054
CollectionRules              : 2
CollectionType               : 0
CollectionVariablesCount     : 0
Comment                     : My first collection created with PowerShell
CurrentStatus                : 5
HasProvisionedMember         : False
IncludeExcludeCollectionsCount : 0
IsBuiltin                    : False
IsReferenceCollection         : False
ISVData                      :
ISVDataSize                  : 0
LastChangeTime               : 5/6/2013 7:39:03 AM
LastMemberChangeTime         : 1/1/1980 12:00:00 AM
LastRefreshTime              : 12/31/1979 11:00:00 PM
LimitingCollectionID         : SMS000001
LimitingCollectionName       : All Systems
LocalMemberCount             : 0
MemberClassName              : SMS_CM_RES_COLL_PS100054
MemberCount                  : 0
MonitoringFlags              :
Name                         : CT All Workstations
OwnedByThisSite               : True
PowerConfigsCount            : 0
RefreshSchedule              : 1
RefreshType                   : True
ReplicateToSubSites           : True
ServiceWindowsCount          : 0

PS PS1:\> New-CMUserCollection -Name "SWU Microsoft Office 2013 Install" -Comment "All Users in this collection will get Microsoft Office 2013 installed" -LimitingCollectionName "All Users"

CollectionID                : PS100055
CollectionRules              : 1
CollectionType               : 0
CollectionVariablesCount     : 0
Comment                     : All Users in this collection will get Microsoft Office 2013 installed
CurrentStatus                : 5
HasProvisionedMember         : False
IncludeExcludeCollectionsCount : 0
IsBuiltin                    : False
IsReferenceCollection         : False
ISVData                      :

```

Figure 9
Creating First
Collections with
PowerShell

Use these commands, respectively:

```
New-CMDeviceCollection -Name "CT All Workstations"
-Comment "My first collection created with PowerShell"
-LimitingCollectionName "All Systems"
```

```
New-CMUserCollection -Name "SWU Microsoft Office 2013
Install" -Comment "All users in this collection will get
Microsoft Office 2013 installed" -LimitingCollectionName
"All Users and User Groups"
```

5. Collections gain their member lists via collection rules. A rule can be a dynamic query or a direct membership rule; it can include or exclude members from another collection. The first of the following sample cmdlets uses PowerShell to create a

direct membership rule that adds an AD group as a member of the SWU Microsoft Office 2013 Install collection. The resource ID can be found by looking at the properties of each object in the Configuration Manager console. The second sample cmdlet shows how to create a dynamic membership query that adds all workstations as members of the CT All Workstation collection.

```
Add-CMUserCollectionDirectMembershipRule -Collectionname
    "SWU Microsoft Office 2013 Install" -ResourceId 2080374411
```

```
Add-CMDeviceCollectionQueryMembershipRule -CollectionName
    "CT All Workstations" -RuleName "All Workstations"
    -QueryExpression 'select * from SMS_R_System where SMS_R_
    System.OperatingSystemNameandVersion like "%workstation%"'
```

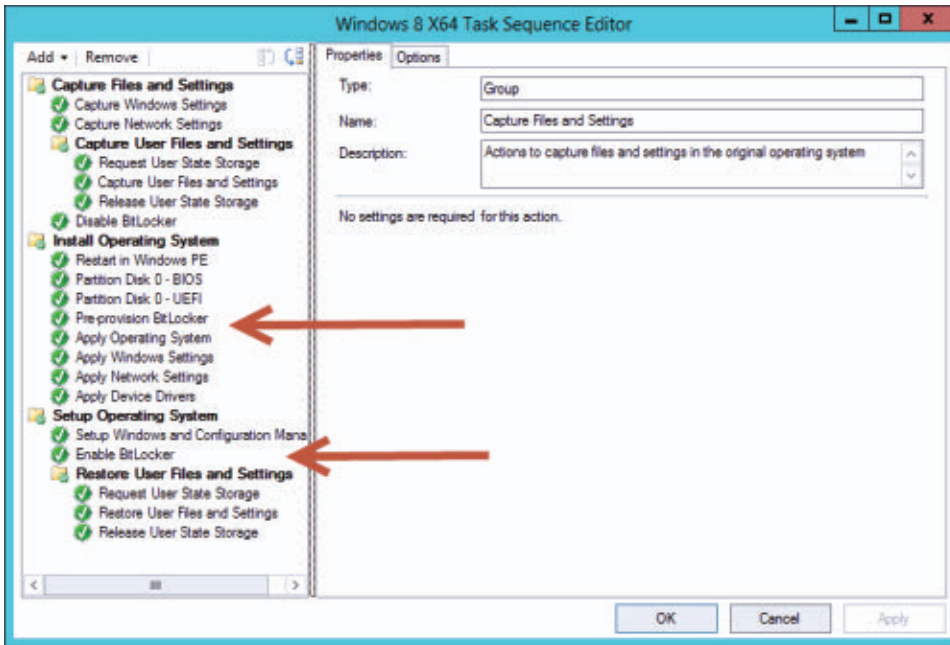
These are just a few examples of how PowerShell comes in handy. With more than 400 cmdlets, the sky is almost the limit of what you can do.

3. New Operating System Deployment Features

In Configuration Manager 2012 release to manufacturing (RTM), Operating System Deployment (OSD) was basically the same as in Configuration Manager 2007. But that has all changed in SP1. Besides finding full Windows 8 and Windows Server 2012 support, you will also find new features that you can use when deploying Windows 7 and even Windows XP computers.

The first thing that you'll notice when upgrading to SP1 is that you need to uninstall the Windows Automated Installation Kit (WAIK) and instead install the Windows Assessment and Deployment Kit (Windows ADK). All your old boot images will be replaced with the standard WinPE 4.0 boot image. These are some of the coolest new features:

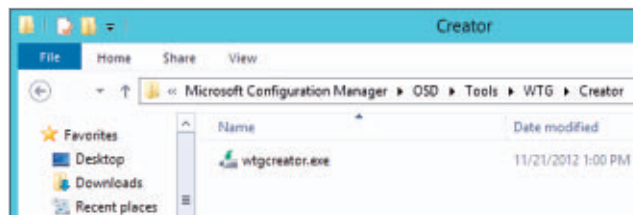
- Windows BitLocker Drive Encryption enhancements allow BitLocker to be provisioned in WinPE (as Figure 10 shows) and

**Figure 10**

New Task Sequence to Preprovision BitLocker in WinPE

encrypt data as it's added. A BitLocker process now takes a few minutes instead of several hours.

- Prestaged media now supports the storage of all content, packages, drivers, and so on. If content changes between media creation and deployment, new content is automatically downloaded from the distribution point. This change is a huge benefit because we often have scenarios in which prestaged media can be several weeks old before reaching its destination. You can also use the prestaged media file and wtgcreator.exe application in the \OSD\Tools\WTG\Creator folder on the site server to create a Windows To Go deployment, as Figure 11 shows.

**Figure 11**

Using wtgcreator.exe with Prestaged Media to Offer a Windows To Go Solution

- Unified Extensible Firmware Interface (UEFI), which replaces the grand old BIOS in newer hardware models, is supported. The main benefits of UEFI are faster boot and support for the latest and greatest hardware.
- New deployment options, as listed in Table 1, provide administrators with much more control of the deployment process.

Table 1: New Deployment Options

Deployment Option	Description
Only Configuration Manager client	This option is useful in refresh scenarios when you want to start the OSD deployment process from a working OS.
Configuration Manager clients, media, and PXE	In this classical deployment option, the task sequence is visible in all environments.
Only media and PXE	This option is used for bare-metal deployments in which the computer isn't booted into an existing OS.
Only media and PXE (hidden)	This option allows the administrator to deploy multiple required task sequences and to automatically select which one to run at deployment. Administrators can use the built-in variable SMSTSPreferredAdvertID.

- Preboot Execution Environment (PXE) provides better logging. Also, the Configuration Manager 2007 monitoring experience is back, which allows administrators to again monitor all the OSD phases in reports and in the Configuration Manager console.
- Much of the support that we're used to seeing in the Microsoft Deployment Toolkit (MDT) is now built into Configuration Manager. Some of the most useful changes add support for additional components in WinPE, such as PowerShell, and add custom files in the boot image.

Configuration Manager SP1 offers you a much better way to control your OSD deployments. OSD is powerful feature, but one failure and you can end up re-imaging 1,000 desktops and 250 servers. Oh yeah, that has happened before—but hopefully won't anymore, thanks to these new options (see Figure 12).

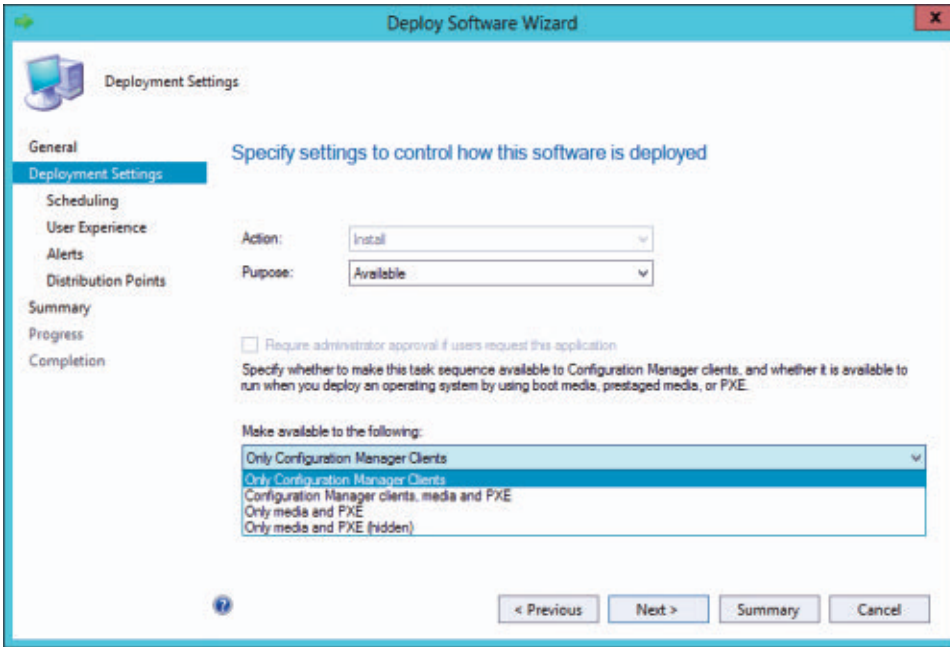


Figure 12
New OSD
Deployments Ensure
More Control

You can use this simple Visual Basic (VB) script in the boot image to select a hidden task sequence:

```
Dim env
set env = CreateObject("Microsoft.SMS.TSEnvironment")
env("SMSTSPreferredAdvertID") = PS100000B
```

4. New Software Update Management Features

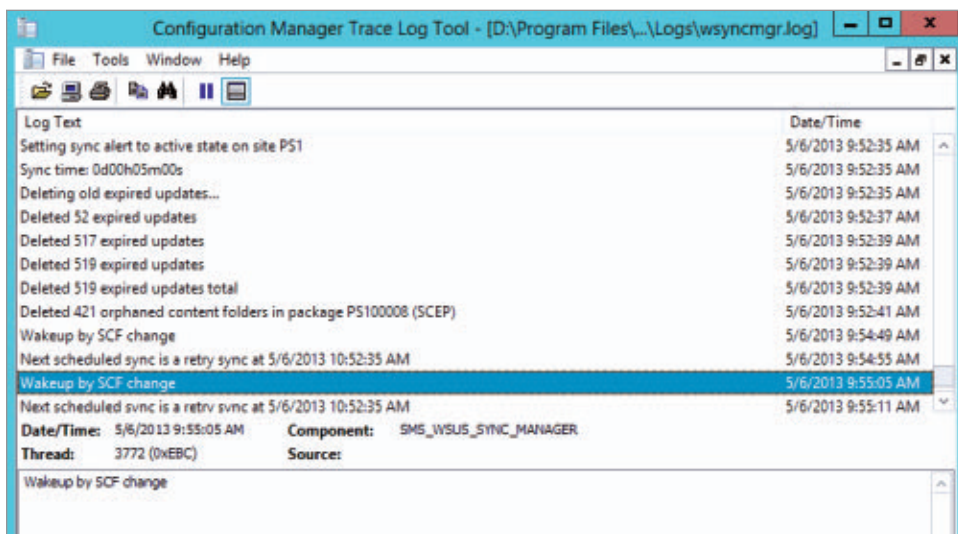
You won't find as many changes in software updates as in some other areas of Configuration Manager. But the changes that you do find can have a huge impact in your environment:

- SP1 includes support for multiple software update points. A limitation in Configuration Manager 2012 RTM was that it supported only one software update point (with the exception of an Internet-based software update point). The change might not sound that big, but it makes a huge difference in environments in which a single primary site covers multiple forests with and

without a trust relationship. Prior to SP1, you needed to allow all clients from all forests and domains access to one software update point—and then deal with the consequences for security and firewalls. Now Configuration Manager supports installation of a software update point, management point, distribution point, and application point (all user-facing site systems) in a remote forest.

- Now you can automatically clean up expired updates from distribution points and source locations, greatly affecting the amount of content that's replicated to distribution points. It isn't uncommon to see software update packages of 10GB to 30GB. Often, 10 to 20 percent of that content is expired; expired updates can't be installed on clients and are a waste of disk space and replication. The cleanup task is fully automated and can't be controlled. You can monitor the cleanup process by reading the `wsyncmgr.log` file on the primary site server, as Figure 13 shows.

Figure 13
Automatic Cleanup
Process



- Allowing fallback to Microsoft Update when updates are unavailable at the distribution point is a new deployment feature that allows Configuration Manager 2012 SP1 clients to fall back to the cloud and download binaries that aren't found locally. This feature

is completely transparent for the end user and isn't the same as allowing the end user real-time access to Microsoft Update.

5. Platform and Infrastructure Changes

Say “cloud,” and many administrators will tell you about many applications—none of which used to be in Configuration Manager. But that all changes with SP1. Now we see cloud integration on site systems, in client support, and—as previously mentioned—as a fallback solution for software updates.

- Windows Intune has long been a standalone cloud-based solution with features such as application deployment, inventory, patch management, and endpoint protection. SP1 introduces a Windows Intune connector that gives the administrator a single pane of glass to manage Windows Intune enrolled devices in the Configuration Manager console. The list of supported devices expands beyond traditional Windows devices (although feature support differs between devices):
 - Apple iOS (iPad and iPhone)
 - Google Android
 - Windows 8 Phone
 - Windows RT
- The cloud-based distribution point is a Windows Azure solution in which content is stored in the cloud. There are several benefits of using a cloud-based distribution point:
 - The cloud distribution point can be used as a fallback solution.
 - Clients will fall back to the cloud distribution point only if the requested content is unavailable on the local or remote distribution point.
 - The cloud-based distribution point can be used by Internet-based clients.
 - The solution doesn't require a full PKI environment.
 - The solution is a dynamic one in which you can change the content requirement on the fly.

- The cloud-based distribution point is managed in the same way as an on-premises distribution point.
- The pull distribution point is a new on-premises distribution point role. A pull distribution point isn't controlled by the site server in the same way. Traffic to the pull distribution point honors neither bandwidth control nor scheduling. Instead, the site server sends a message to the pull distribution point, informing it that content is available and can be downloaded from one of the pull distribution point partners.
- Another infrastructure change is the ability to add a new central administration site to an existing primary site. This can be done only once in the hierarchy.
- Migration from other Configuration Manager 2012 SP1 sites is also supported. Previously, migration was supported only from a Configuration Manager 2007 SP2 environment.
- Microsoft SQL Server replication has been optimized, which is extremely useful when you're working with multiple sites.
- There's support for Mac OS clients. The supported features are application deployment, settings management, and inventory management. Mac OS support requires that you implement PKI and have HTTP Secure (HTTPS) support on at least one management point, distribution point, enrollment point, and enrollment proxy point.
- There's support for UNIX and Linux servers. Supported features include malware protection, software deployment, and inventory management. The UNIX support doesn't require any changes in the infrastructure.

More Than Expected

I hope I've proven that Configuration Manager 2012 SP1 is more than just another service pack with a few bug fixes. With support for the cloud and new OSs (Microsoft and others), as well as several feature improvements, this pack offers much more than you might expect. ■

Windows Server 2012 Storage Live Migration

Become a migration master

Windows Server 2012 brought new levels of mobility to the virtual environment. This mobility extends beyond the previous live migration capability, which was limited to migration within a cluster with shared storage. Windows Server 2012 introduces migration of virtual machines (VMs) between any Windows Server 2012 Hyper-V hosts, standalone or clustered. The cluster is no longer a mobility boundary, so enterprises have complete flexibility. Often the requirement is to move not the VM but rather its storage, something that was possible prior to Windows Server 2012 only after shutting down the VM.

Windows Server 2012 supports three main types of storage for VMs: DAS; SAN-based (typically connected via Fibre Channel or iSCSI); and—new to Windows Server 2012—support for Server Message Block (SMB) 3.0 file shares, such as those hosted on a Windows Server 2012 file server or any NAS/SAN that has SMB 3.0 support. Windows Server 2012 storage live migration allows the storage used by a VM, including the VM's configuration and virtual hard disks (VHDs), to be moved between any supported storage, with zero downtime to the VM. Migration to a different folder on the same disk, between LUNs on the same SAN, from DAS to SAN, from SAN to an SMB file share—if the storage is supported by Hyper-V, then VMs can be moved with no downtime. Note that storage live migration can't move non-virtualized storage, so if a VM is using pass-through storage, then it can't be moved. The good news is that with the new VHDX format (which allows 64TB VHDs), there's no reason to use pass-through storage, from either a size or performance perspective.



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley).



Email



Twitter



Website

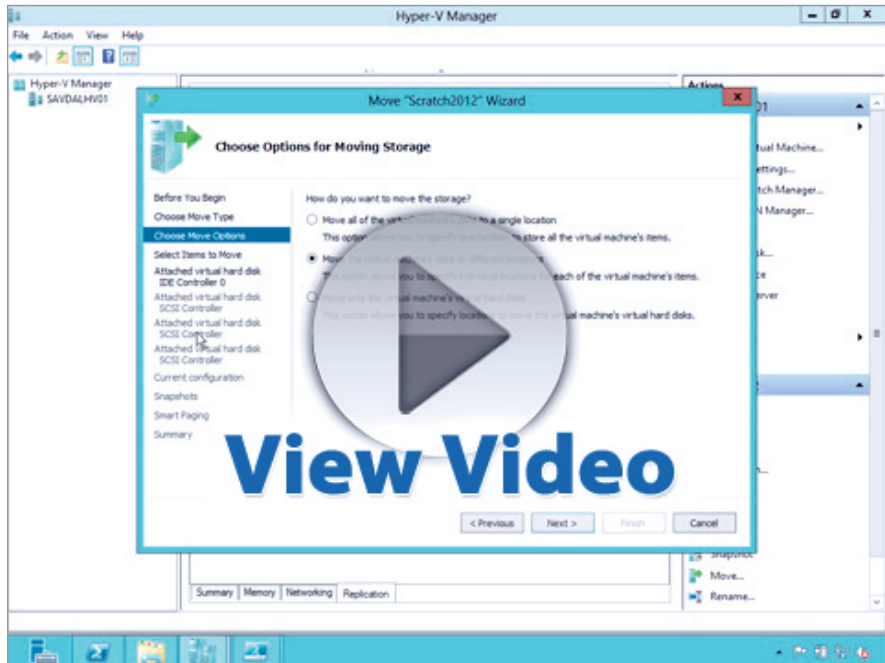


Blog

Video



John Savill
demonstrates
Windows Server 2012's
Storage Migration
feature



The ability to move the storage of a VM at any time, without affecting the availability of the VM, is vital in two key scenarios:

- The organization acquires new storage, such as a new SAN, or is migrating to a new SMB 3.0 appliance and needs to move VMs as part of a planned migration effort.
- The storage in the environment is out of space or can't keep up with the I/O operations per second (IOPS) requirements, and VMs need to be moved as a matter of urgency. In my experience, this scenario is the most common.

How Storage Live Migration Works

The mechanics behind Windows Server 2012 storage live migration are quite simple but provide the most optimal migration process. Remember that the *VM* isn't moving between hosts (although you can use shared-nothing live migration to accomplish that); only the *storage* moves from a source location to a target location.

Storage live migration uses a one-pass copy of VHDs. The pass works as follows:

1. Storage live migration is initiated from the GUI or Windows PowerShell.
2. The copy of the source VHDs, smart paging file, snapshots, and configuration files to the target location is initiated.
3. When the copy initiates, all writes are performed on the source and target VHD through a mirroring process in the virtual storage stack.
4. After the copy of the VHDs is complete, the VM is switched to use the VHDs on the target location. (The target is up-to-date because all writes are mirrored to the target while the copy is in progress.)
5. The VHDs and configuration files are deleted from the source.

The actual storage live migration process is managed by the Virtual Machine Management Service (VMMS) in the parent partition. However, the heavy lifting of storage live migration is performed by the VMs' worker process and the storage virtualization service provider in the parent partition. The mechanism for the storage copy is just an unbuffered copy operation plus the additional I/O on the target for the mirroring of writes during the copy. In reality, the additional I/O for the ongoing writes is negligible compared with the main unbuffered file copy. The path used is whichever path exists to the target: iSCSI or Fibre Channel for a SAN target, whichever network adapter or adapters have a path to the share for SMB. Any underlying storage technologies that optimize performance are fully utilized. If you're copying to or from SMB and using NIC Teaming, SMB Direct, or SMB Multichannel, then those technologies will be used. If you're using a SAN that supports offloaded data transfer (ODX) and you're moving a VM within a LUN or between LUNs, then ODX will be used, meaning that the move will use almost no load on the host and will complete very quickly.

The SAN ODX scenario is the best case. For all other situations, it's important to realize exactly what an unbuffered copy means to

your system. The unbuffered copy is used because during storage live migration, you don't want to use a large amount of system memory for caching of data on a virtualization host.

Performing a copy can cause a significant amount of I/O load on your system for both reading the source and writing to the target. To get an idea, try manually creating an unbuffered copy on your system by using the Xcopy command with the /J switch. This creates a similar load to what a storage live migration would inflict on your system, again considering that the ongoing mirrored writes are negligible. Therefore, consider moving a VM between folders on a local disk (likely to be a worst-case scenario). The data would be read from and written to the same disk, causing a huge amount of disk thrashing; it would likely take a long time and would adversely affect any other VMs that use that disk. If the source and target are different storage devices, then the additional load won't be as severe as a local move—but must still be considered.

Moving a VM causes nothing Hyper-V-specific about the disk I/O, which is the same as for any data-migration technology (although other technologies might not have capabilities such as ODX when a SAN is involved). Ultimately, the data must be read and written. This doesn't mean that you shouldn't use storage live migration, but it does mean that you should plan carefully when you use it.

You probably won't want to perform the migration during normal working hours because of the possible adverse effect to other loads. I suspect this is why no automated storage live migration process is part of the Dynamic Optimization in System Center Virtual Machine Manager (VMM) 2012, which rebalances VMs within a cluster. If you detect a large I/O load on a storage subsystem in the middle of a weekday, the last thing you want to do is add a huge extra load by trying to move things around. The best option is to track I/O over time, then move the VM's storage at a quiet time—a task that's easy to script with PowerShell or to automate with technologies such as Microsoft System Center Orchestrator 2012.

Configuring Storage Live Migration

If you've installed the Hyper-V role on your server, you're all done. No specific configuration is needed to use storage live migration; it just works. As previously stated, storage live migration uses whichever path exists to communicate with the source and target storage, and it's enabled by default (in fact, you can't disable it). The only configuration is that you can set how many simultaneous storage live migrations are allowed. To do so, use the Hyper-V Settings action. In the Storage Migrations area, set the desired Simultaneous storage migrations number, as Figure 1 shows.

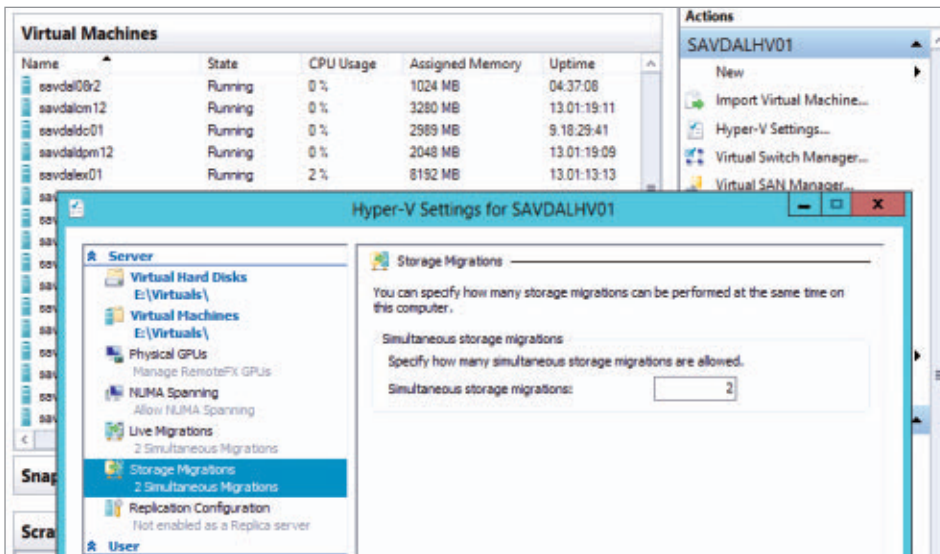


Figure 1
Setting the Number of
Simultaneous Storage
Live Migrations

You can also configure this setting by using PowerShell:

```
Set-VMHost -MaximumStorageMigrations <number to allow>
```

You only need extra configuration if you're using SMB storage for the migration target and are initiating the migration remotely, either through Hyper-V Manager or PowerShell. In other words, you aren't running the tools on the actual Hyper-V host. This type of remote management

is preferred for Windows Server 2012; all management should be performed remotely, using PowerShell or from a [Windows 8](#) machine.

When you configure SMB storage for use with Hyper-V, you need to set several specific permissions, including giving administrators full control to create a VM on SMB or to move to SMB as part of a storage live migration, as their credential is used. As I explain in the article “[Shared-Nothing VM Live Migration with Windows Server 2012 Hyper-V](#),” remotely initiating a shared-nothing live migration requires the configuration of Kerberos constrained delegation on each Hyper-V server. The Microsoft Virtual System Migration Service requires this configuration because by default, a Windows server can’t pass a credential that’s being used on the server to another server. Doing so would generally be bad from a security perspective but is exactly what we need here and is acceptable in this specific, scoped context:

1. The administrator initiates the storage live migration remotely through Hyper-V Manager or PowerShell remoting. The administrator’s current credential is passed to the host that’s performing the action, or a specific credential may be passed, if you’re using PowerShell.
2. The server performing the storage live migration must then connect to the SMB share and create files. To do so, it needs to use the administrator’s credential. However, doing so would be passing on the credential (aka delegation), which isn’t allowed by default.

To enable this scenario, you must enable Common Internet File System (CIFS) constrained delegation for each Hyper-V server to the various SMB file servers. This task is a simple one:

1. Launch Active Directory Users and Computers.
2. Navigate to your Hyper-V servers, right-click one, and choose Properties.
3. Choose the Delegation tab.
4. Make sure that the *Trust this computer for delegation to specified services only* and *Use Kerberos only* options are selected.

5. Click Add.
6. Click *Users or Computers*, choose your SMB file servers, and click OK.
7. In the list of available services, select *cifs* for each server, and click OK, as Figure 2 shows.

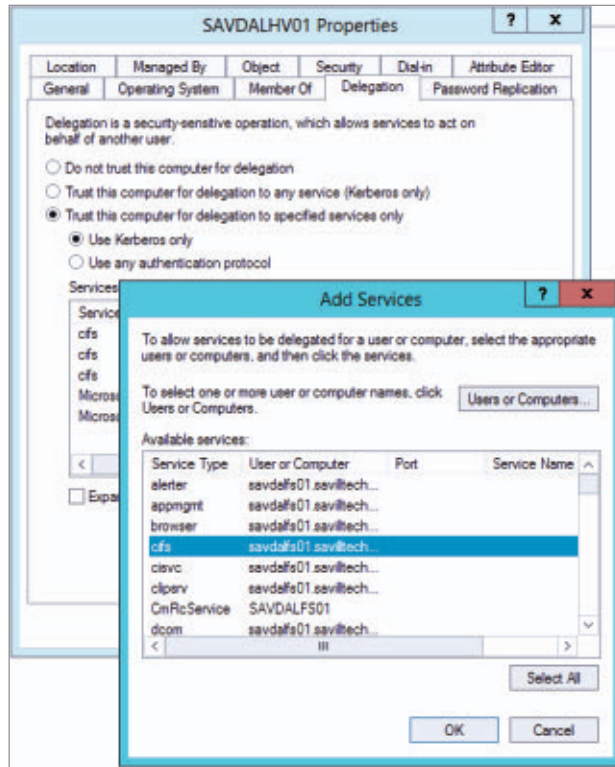


Figure 2
Enabling Kerberos
Constrained
Delegation to the File
Servers for CIFS

You can now remotely trigger storage live migrations, even to SMB storage.

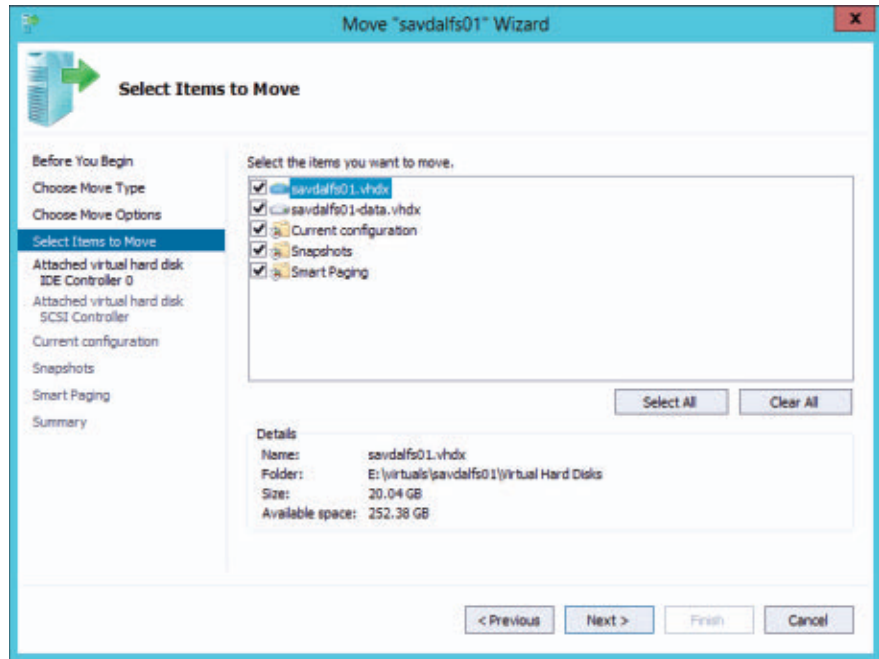
Performing Storage Live Migration

Now that the environment is ready for storage live migrations, all that's left is to perform them. Storage live migrations can be triggered through Hyper-V Manager or through PowerShell. You have two options when performing a storage live migration. You can move everything to one location, or you can choose different locations for each item that's

stored as part of a VM (i.e., one location for the configuration file, one for the snapshots, one for smart paging, one for VHD1, one for VHD2, and so on), as Figure 3 shows. This isn't a problem when using graphical tools but adds an interesting aspect when using PowerShell.

Figure 3

Selecting Items to Move



Start by using Hyper-V Manager to perform the move. Doing so helps you understand the available options:

1. Launch Hyper-V Manager.
2. Choose the VM with the storage that needs to be moved and choose the Move action.
3. Click Next to proceed to the Before You Begin page of the wizard.
4. Choose the *Move the virtual machine's storage* option (since you're only moving the storage).
5. You can now choose to move all the VM's data to a single location, which is the default, or to move the data to different locations, or to move only the VHDs for the VM. Make your selection and click Next.

6. If you chose the default, you're prompted for the new storage location; specify it, and then click Next. If you chose either of the other options, you're shown a separate page on which you must select the target location for each element of the VM's data; set the location for each item, and then click Next.
7. Review your options and click Finish to initiate the storage live migration.

To perform the storage live migration from PowerShell, use the `Move-VMStorage` cmdlet. If you're moving everything to a single location, simply pass the VM name and the new target location with the `DestinationStoragePath` parameter. (Note that a subfolder with the VM name isn't created automatically. If you want the VM in its own subfolder, you need to specify that as part of the target path.) Here's an example:

```
Move-VMStorage -DestinationStoragePath <target path> -VMName
<vmname>
```

If you want to move separate data to different locations, the process is more complicated. Instead of using `DestinationStoragePath`, use the `SmartPagingFilePath`, `SnapshotFilePath`, and `VirtualMachinePath` parameters to pass the location for the smart paging file, snapshots, and VM configuration, respectively. For the VHDs, use the `Vhds` parameter. However, you can have more than one VHD per VM—in fact, you can have hundreds of them—and PowerShell doesn't like an arbitrary number of parameters. Therefore, to pass the VHDs' new location, you need to create a hash value for the `SourceFilePath` and `DestinationFilePath` for each VHD, and then place them into an array, which is passed to the `Vhds` parameter. Pleasant!

The following example moves a VM with three VHDs, a smart paging file, configuration, and snapshots. Note you don't need to move all elements of a VM; you only need to specify the pieces that you want to move. Other unspecified elements stay in their current

location. Note that in the example, the hash values (value pairs) use curly brackets { } whereas the array uses parentheses ().

```
Move-VMStorage -VMName <vmname> -SmartPagingFilePath d<smart paging
file path> -SnapshotFilePath <snapshot path> -VirtualMachinePath
<vm configuration path> -Vhds @(@{ "SourceFilePath " = "C:\vm\
vhd1.vhdx "; "DestinationFilePath " = "D:\VHDs\vhd1.vhdx "}, @{
"SourceFilePath " = "C:\vm\vhd2.vhdx "; "DestinationFilePath " =
"E:\VHDs\vhd2.vhdx "}, @{ "SourceFilePath " = "C:\vm\vhd3.vhdx ";
"DestinationFilePath " = "F:\VHDs\vhd3.vhdx "})
```

When the storage live migration is initiated, it runs until it's finished, no matter how long that might take. As the administrator, you can cancel the storage live migration manually by using the *Cancel move storage* action. Rebooting the Hyper-V host also cancels all storage live migrations. You can see the progress of storage live migrations in the Hyper-V Manager tool or by querying them through Windows Management Instrumentation (WMI):

```
PS C:\ > Get-WmiObject -Namespace root\virtualization\v2 -Class Msvm_
MigrationJob | ft Name, JobStatus, PercentComplete, VirtualSystemName
```

Name	JobStatus	PercentComplete	VirtualSystemName
----	-----	-----	-----
Moving Storage	Job is running	14	6A7C0DEF-9805-4242-92F9-98E6F...

Migrate Responsibly

Storage live migration is a great new capability for Hyper-V, if you use it wisely. The feature gives organizations new flexibility in implementing new storage without affecting the availability of services. You can even use it to rebalance storage subsystems with uneven loading—but be sure to plan your migrations to minimize I/O impact. ■

FAQ

Answers to Your Questions

Q: How do I force the Start screen to always show on my main display in Windows 8.1?

A: Windows 8.1 allows the primary monitor to be used to display the Start screen. To make the change, perform the following:

1. Right-click the taskbar and select Properties.
2. Select the Navigation tab.
3. Under the Start screen area, check the option *Always show Start on my main display when I press the Windows logo key*.
4. Click OK.

—John Savill



John Savill

Q: If I use pass-through storage with Windows Server 2012 Hyper-V, what features do I lose?

A: Pass-through storage is a configuration that lets a **virtual machine** (VM) directly access a disk on the Hyper-V host. In this configuration, the host must have the disk in an offline state, and only the VM can access the disk in pass-through configuration, making it an exclusive resource.

Pass-through storage was required in previous versions of Hyper-V due to limitations in the virtual hard disk (VHD) format (mostly because its maximum size allowable was 4TB). **Windows Server 2012**, however, uses the new VHDX format.

This new VHDX format not only has a new maximum size of 64TB but also delivers performance-matching native disk levels, even with



Jan De Clercq

dynamic disks, removing the need for pass-through. If you do use a pass-through disk, you lose such key features or abilities as these:

- Create a snapshot
- VM backup
- Storage Migration
- Hyper-V Replica
- Storage quality of service (QoS)

—John Savill

Q: What's the easiest way to create a global audit policy that automatically logs events for all administrator changes to the system registry on all the domain controllers (DCs) in a Windows domain?

A: To set up a global audit policy, you can leverage a Windows feature called Global Object Access Auditing, which Microsoft introduced in [Windows Server 2008 R2](#). A global object access audit policy can be used to enforce an object access audit policy for a file system or registry folder, without having to configure and propagate conventional system ACL (SACL) settings on each machine. You can find a good introduction to this feature on TechNet's [Global Object Access Auditing](#) page.

To configure, apply, and validate a global object access audit policy for administrator changes to the system registry on your DCs, follow these steps:

1. Log on to your domain as a member of the local Administrators group and start the Group Policy Management Console (GPMC).
2. In the console tree, navigate to \Domains\< Your_ Domain > \Group Policy Objects\Default Domain Controllers Policy, where < Your_ Domain > is the name of your domain. Right-click Default Domain Controllers Policy and click Edit.
3. In the Group Policy Management Editor, navigate to the \Computer Configuration\Policies\Windows Settings\Security

Settings\Advanced Audit Policy Configuration\System Audit Policies container.

4. Double-click Object Access, then double-click Audit Registry. Select the *Configure the following audit events* check box, select the Success and Failure check boxes, and click OK.
5. Double-click Global Object Access Policies, then double-click Registry. Select the *Define this policy setting* check box and click Configure.
6. In the Advanced Security Settings for Global Registry SACL box, click Add. Add all default administrator groups (e.g., Domain Admins, Enterprise Admins) to the list and other custom administrator groups that you've defined and want to audit.
7. In the Auditing Entry for Global Registry SACL box, select the Successful or Failed activities (e.g., Create Subkey, Delete, Change Permissions, Read) for which you want to log audit entries.
8. Click OK three times to complete the audit policy configuration.
9. Apply the Group Policy Object (GPO) change. On each of your DCs, open a command prompt and run this command:

```
gpupdate /force
```

—Jan De Clercq

Q: Is it better to have one large 64-node cluster in Windows Server 2012 or multiple, smaller clusters?

A: There's no right or wrong answer to this question. The nodes in a cluster have full ability to share resources and to move resources between the nodes. Thus, the larger the number of nodes in a cluster, the greater the ability to consolidate and share resources.

This leads to less waste and allows the number of “spare” hosts to be reduced. Additionally, [Windows Server 2012](#) has capabilities such

as one-click updating of an entire cluster, so the smaller the number of clusters, the lower the amount of management overhead.

The opposite argument is that a cluster is still a single cluster configuration, so in the event of some very severe cluster corruption, all the nodes in a cluster would be unavailable. Because of this, typically people will at least split resources into two clusters; remember, in Windows Server 2012, it's still possible to move virtual machines (VMs) between clusters without downtime.

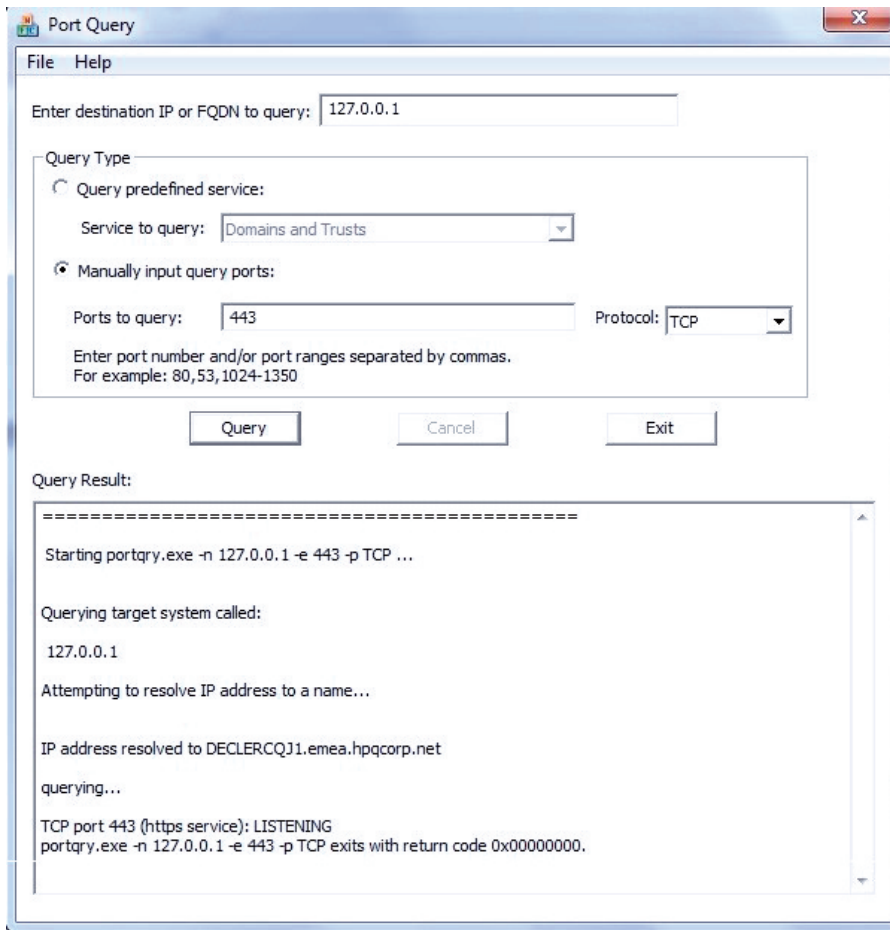
—John Savill

Q: What GUI tool can I use to check the status of the TCP and UDP ports on my local computer or on a remote computer?

A: A nice tool that you can use is Microsoft's free [PortQueryUI](#). It provides a GUI for the portqry.exe command-line tool, which is installed by default on recent Windows OS versions. To use it, for example, to query the status of TCP port 443 on your local computer using PortQueryUI, follow these steps:

1. Download and install the PortQueryUI tool on your local computer. (The download includes the portqry.exe command-line tool, in case you're running an earlier Windows OS.)
2. Go to the C:\PortQryUI folder on your local system and run portqueryui.exe.
3. In the Port Query dialog box, enter the destination IP address or Fully Qualified Domain Name (FQDN) to query. In this case, you can enter 127.0.0.1 or localhost.
4. Select *Manually input query ports* and type 443 in the *Ports to query* field. (By default, TCP is selected in the Protocol drop-down list.)
5. Click the Query button. The status of port 443 is displayed in the bottom pane of the Port Query dialog box, as Figure 1 shows.

Figure 1
Using PortQueryUI
to Display the Status
of Port 443



—Jan De Clercq

Q: I'm using SharePoint 2013 in a basic lab environment but pages are taking a long time to load—how can I speed it up?

A: SharePoint 2013 has fairly high performance requirements, so make sure you meet these. I've seen a lack of memory cause problems. SharePoint 2013 now uses FAST technology for its search, but if your SharePoint site isn't using search, you could

disable this service to free up resources. In Services, stop the following and set to start as Manual:

- SharePoint Search Host Controller
- SharePoint Server Search 15

This might help speed up the response of the SharePoint server. However, realize that this is going to stop Search results updating, which is typically a key benefit of SharePoint.

—John Savill

Q: What has changed in the quorum model in Windows Server 2012?

A: At times, the quorum model in [Windows Server 2012](#) could be difficult to understand and architect, such as knowing when to use a disk or file share witness, and handling situations when nodes were taken offline, which then changed the need for an additional witness. In Windows Server 2012 R2, this is now much simpler. The guidance in Windows Server 2012 R2 is to always configure a witness. If you've got shared storage, use a disk witness; if there's no shared storage, use a file share witness. The Windows Server 2012 R2 failover cluster will use the witness only if it needs it.

For example, if there are an even number of nodes, then the witness will be used and will be given a vote. If there are an odd number of nodes, then the witness isn't used and won't have a vote. This distribution and vote giving is all done automatically. It really simplifies quorum configurations and was a feature that was needed for working with the dynamic quorum capability added in Windows Server 2012. ■

—John Savill

Product News for IT Pros

Imation Touts IronKey at TechEd

At Microsoft TechEd 2013, we tinkered with Imation's impressive array of IronKey encrypted flash and external hard disk drives, and we also looked at the company's recent IronKey Workspace, which lets organizations outfit mobile professionals with a secure, fast USB platform to run Windows To Go from a USB stick on multiple compatible PCs—an ideal solution for teleworkers, contractors, and those implementing BYOD strategies.

Many organizations are seeking smart, easy-to-use, and highly secure solutions for employees to safely transport their private data, intellectual property, and brands. The IronKey D80 flash drives and H80 hard drives feature tamper-resistant enclosures and hardware-based 256-bit AES encryption with strong authentication to reliably keep data private and secure. IronKey D80 and H80 drives are ideal choices for organizations seeking compliance with most business policies and standards and that recognize the superior security offered by hardware encryption. For more information about Imation's IronKey D80 flash drives and H80 hard drives, check out the [Imation Mobile Security Store](#).

The IronKey Workspace is certified for deployment of Windows To Go, a [Windows 8](#) feature that lets enterprise users boot a full version of Windows 8 from an external USB drive on compatible host PCs. Organizations running Windows 8 Enterprise Edition can provision a Windows desktop onto the IronKey Workspace to create a “PC on a Stick”—with the OS and data contained on the USB drive—while still leveraging the host PC's hardware and resources, such as monitors, cameras, and network connections. For more information, visit the [Imation website](#).





Dell Announces Migration, Management, and Monitoring of Microsoft Environments

Dell announced a series of enhancements designed to help organizations optimize the migration, management, and monitoring of Microsoft environments. The company has expanded its portfolio of systems and information management solutions to ensure optimal visibility, availability, and performance of the IT infrastructure. New releases include the latest version of Spotlight for SQL Server Enterprise, Spotlight Project Lucy (a new cloud-based set of productivity tools that let SQL Server users obtain a free system health check), enhancements to Dell KACE K2000 deployment appliances, and a new release of MessageStats Business Insights, which extends Microsoft's Exchange Server 2013 data loss prevention (DLP) reporting to organizations using on-premises Exchange Server 2013. For more information, check out the [Dell website](#).



IS Decisions' UserLock 7.0 Introduces Wi-Fi Session Control

IS Decisions announced the availability of UserLock 7.0, its identity and access management solution. The highlight of UserLock 7.0 is its Wi-Fi session control capability, which enables organizations to mitigate the increasing network threats that accompany today's growing BYOD trend. By limiting concurrent logons and restricting user logons according to customized access policies, UserLock helps IT control insider threats and prevent password sharing. "In today's workplace, BYOD is quickly becoming the rule rather than the exception, but native Windows Server functionality does not provide adequate means to control user access from personal devices," said Francois Amigorena, president and CEO of IS Decisions. "We've designed UserLock 7.0 to alleviate this increased risk to corporate security by empowering IT to track, record, and automatically block all inappropriate or suspicious sessions, including Wi-Fi or IIS." For more information, visit the [IS Decisions website](#).

OCC's Procyon Offers Complete System Approach to Structured Cabling for Data Centers



Optical Cable Corporation (OCC) introduced Procyon, a new family of structured cabling products that offers a complete system approach to data center design and connectivity. Procyon integrates essential accessibility and cable management features, along with the highest density per rack unit on the market. Other options might have high density but lack effective cable management, or they might have cable management covered but not the density or available port count. The Procyon family of products includes data center cabinets, copper and fiber panels with integrated cable management systems, high-density fiber cassettes, and more. The Procyon systems are designed for pre-configuration so that IT pros can kit together the components needed to simplify installations. For additional information about the Procyon family of data center solutions, visit the [OCC website](#).

A10 Networks Expands Thunder Series



A10 Networks announced its new entry-level A10 Thunder Series, extending its family of Unified Application Service Gateways (UASGs) to small-to-mid-sized businesses (SMBs) and enterprise customers. Joining the high-end Thunder 6430(S) and 5430S are the new Thunder 3030S, 1030S, and 930, which expand the line of next-generation Application Delivery Controllers (ADCs). All are compact 1 rack-unit (RU) appliances that provide advanced functionality by consolidating premium, all-inclusive application services in a single and easy-to-manage platform. “Our new Thunder models extend the expansive benefits of our UASG family to the entry-level and mid-range markets,” said Lee Chen, founder and CEO of A10 Networks. “With the new Thunder 3030S, 1030S, and 930 UASGs, customers receive additional value to optimize and scale their existing infrastructure, reduce latency and cost through device consolidation, and increase management efficiency.” For more information, visit the [A10 Networks website](#).



Paessler Introduces New Passive Application Performance Sensor

Paessler introduced a new monitoring technology in the latest Stable version of PRTG, which lets users monitor networks 24 × 7. The new PassiveApplication Performance Sensor lets you monitor the performance of a server or service without accessing either the client or the server directly. For most monitoring scenarios, you either need access to the server or device itself (to monitor vital data such as CPU, memory, disks, bandwidth) or access to the service (to send simulated requests to the server and look at the timing and the content of the replies). The new sensor type uses a completely different approach: This sensor applies PRTG's built-in packet sniffer to look at all TCP packets going into a server, and it checks the reply packets from the server. The idea is that if you measure the time between a TCP packet roundtrip, you can measure the performance of the service or server. For more information, visit the [Paessler website](#).



Asigra Announces Cloud-to-Cloud Backup Support

Asigra announced cloud-to-cloud backup capability for Google Apps, supporting multiple tier-one cloud applications and platforms, including Salesforce, IBM SmartCloud, and Google Apps. Now included in Asigra Cloud Backup 12.2, this new functionality ensures data recovery while ending the requirement for multiple application/platform-centric backup solutions to protect physical, virtual, cloud, and mobile computing platforms. Users can automatically protect Google Apps data at all times. The latest advancement allows companies to back up and restore all important business information in Google Apps, including email messages, calendars, contacts, documents, and sites. Users can automate and schedule the backup activities for the data in Google Apps, select the number of generations of the information that need protection, set retention rules, and even determine separate backup frequencies for different sets of data. Visit the [Asigra website](#) to learn more. ■

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowssitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
forums.windowssitpro.com

News

Check out the current news and information about Microsoft Windows technologies.
www.winsupersite.com

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
windowssitpro.com/vip-premium-membership

SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

IT/Dev Connections..... 61

Windows IT Pro..... 1, 2, 31, 32, 45

Vendor Directory

A10 Networks 105

Apple 85

Asigra 106

Dell..... 104

Google 85

Imation 103

IS Decisions 104

Openfiler..... 48

Optical Cable 105

Paessler..... 106

StarWind Software..... 48

VMware 47

WindowsITPro